CHAPTER 4: FOLIATIONS AND FLOER THEORIES

DANNY CALEGARI

Abstract. These are notes on the theory of taut foliations on 3-manifolds, which are being transformed into Chapter 4 of a book on 3-Manifolds. These notes follow a course given at the University of Chicago in Spring 2016.

Contents

1. Foliations

In this section we collect basic facts and definitions about foliations in general, specializing to some extent to codimension one foliations, but not yet on 3-manifolds.

1.1. **The definition of a foliation.** For any n and $p \le n$ we can fill up $\mathbb{R}^n$ with parallel copies of $\mathbb{R}^p$ (e.g. the subspaces for which the last $n-p$ coordinates are constant). We call this the *product foliation* of $\mathbb{R}^n$ by coordinate $\mathbb{R}^p$'s.

A *codimension* q foliation on an n-manifold is a structure locally modeled on the product foliation of $\mathbb{R}^n$ by coordinate $\mathbb{R}^p$'s, where $p = n - q$. Formally, it is a decomposition of the manifold into disjoint embedded p-submanifolds (called *leaves*) so that locally (i.e. when restricted to sufficiently small open sets around any given point) the components of the leaves partition the open set in the same way the product foliation of $\mathbb{R}^n$ is partitioned into coordinate $\mathbb{R}^p$'s.

The issue of smoothness is important for foliations. We say a foliation is *leafwise* C^r if each leaf is a C^r submanifold. Providing $r \ge 1$, the tangent spaces to the leaves give rise to a p-dimensional subbundle $T\mathcal{F}$ of TM. We say a foliation is C^r if $\mathcal{F}$ is leafwise C^r, and if $T\mathcal{F}$ is a C^r subbundle. We will occasionally consider foliations which are no more regular

Date: January 23, 2020.

than C^1 or C^0, but we will typically insist that they are leafwise C^∞ (this turns out not to be a restriction in low dimensions).

Depending on the degree of smoothness, the data of a foliation may be given in several ways, which we now discuss.

1.1.1. *Involutive distributions.*

Definition 1.1. A (smooth) p-dimensional subbundle ξ of TM is *involutive* (one also says *integrable*) if $\Gamma(\xi)$ is a Lie algebra; i.e. if, whenever X, Y are vector fields tangent to ξ, so is $[X, Y]$.

Frobenius' Theorem says that a bundle is involutive if and only if there is a smooth p-dimensional submanifold passing through each point of M and everywhere tangent to ξ. The (germ of such a) manifold is unique if ξ is smooth, and the (maximal) submanifolds passing through different points are disjoint or equal, and are precisely the leaves of a foliation of M.

Thus a smooth subbundle ξ of TM is equal to $T\mathcal{F}$ for some smooth foliation $\mathcal{F}$ if and only if it is involutive.

1.1.2. *Differential ideals.*

Definition 1.2. Let ξ be a p-dimensional subbundle of TM. A form ω *annihilates* ξ if $\omega(X_1, \cdots, X_p) = 0$ pointwise for all sections $X_i \in \Gamma(\xi)$.

The set of forms $I(\xi)$ annihilating ξ is an ideal in $\Omega^*(M)$; i.e. it is closed under wedge product. Furthermore, an ideal in $\Omega^*(M)$ is of the form $I(\xi)$ for some ξ as above if and only if it is locally generated (as an ideal) by $n-p$ independent 1-forms.

An ideal I in $\Omega^*(M)$ is said to be a *differential ideal* if it is closed under exterior derivative d. There is a duality between differential ideals and Lie algebras; under this duality, Frobenius Theorem becomes the proposition that an ideal $I(\xi)$ is a differential ideal if and only if ξ is involutive.

Example 1.3 (Foliations from 1-forms). Let M be an n-manifold. Suppose on some open $U \subset M$ we have $q = n - p$ independent 1-forms $\omega_1, \cdots, \omega_p$. Let ξ be the kernel of the ω_i (i.e. the p-dimensional subbundle of TU where all ω_i vanish). Then ξ is involutive if and only if there are 1-forms α_{ij} so that

$$d\omega_i = \sum_j \alpha_{ij} \wedge \omega_j$$

for all i.

If M is a 3-manifold, and ξ is a 2-dimensional distribution, then $\xi = T\mathcal{F}$ for some $\mathcal{F}$ is and only if locally $\xi = \ker(\omega)$ for some nonzero 1-form ω with $d\omega = \alpha \wedge \omega$. Equivalently, $\omega \wedge d\omega = 0$.

1.1.3. *Charts.* Another way to define the structure of a foliation is with charts and transition functions. On $\mathbb{R}^n$ let x denote the first p coordinates, and y the last q coordinates, where $p + q = n$. Then the data of a foliation on M is given by an open cover of M

by charts U_α and homeomorphisms $\varphi_\alpha : U_\alpha \to \mathbb{R}^n$ so that on the overlaps the transition functions $\phi_{\alpha\beta} := \varphi_\beta \varphi_\alpha^{-1}$ have the form

$$\phi_{\alpha\beta}(x, y) = (\phi^x_{\alpha\beta}(x, y), \phi^y_{\alpha\beta}(y))$$

In other words, the y coordinate of $\phi_{\alpha\beta}(x, y)$ only depends on y, and not on x.

Transition functions of this form take leaves of the product foliation into other leaves. Thus we can pull back leaves under the charts to define the leaves of a foliation on M. Notice that this definition makes sense for foliations which are only C^0.

Example 1.4. The most basic examples of foliations are the product foliations of $\mathbb{R}^n$. More generally, if M, N are manifolds of dimension p and q respectively, then $M \times N$ has a p-dimensional foliation by factors $M \times$ point and a q-dimensional foliation by factors point $\times$ N.

Example 1.5. If $\mathcal{F}$ is a foliation of M, then the restriction of $\mathcal{F}$ to any open submanifold U of M is also a foliation.

Example 1.6. If M is any manifold, and X is a nonsingular vector field then TX is evidently involutive, and is equal to $T\mathcal{F}$ for some 1-dimensional foliation $\mathcal{F}$. The leaves of $\mathcal{F}$ are simply the *integral curves* of the flow associated to X.

1.2. **Transversals and Holonomy.** Let $\mathcal{F}$ be a foliation of dimension p on an n-manifold M. A *transversal* τ is an (open) submanifold of dimension q which intersects the leaves of $\mathcal{F}$ transversely; i.e. in a local product chart in which the (components of the) leaves are the submanifolds of $\mathbb{R}^n$ with last q coordinates constant, τ is the submanifold with first p coordinates equal to zero.

If M is closed, we may find a finite set of transversals $\tau_1, \cdots, \tau_k$ which together intersect every leaf of $\mathcal{F}$. Denote by τ the union of the τ_i.

Let γ be a path contained in a leaf λ of $\mathcal{F}$ which starts and ends on τ. We can cover γ by (finitely many) product charts. At $\gamma(0) \subset \tau$ we can use the y coordinate of the chart as a parameterization of the germ of τ near $\gamma(0)$. When we move from one chart to the next, the form of the transition function shows that there is a well-defined map on y-coordinates, so we get the germ of a map between open subsets of $\mathbb{R}^q$ associated to each transition. The composition of the finitely many transitions as we move from $\gamma(0)$ to $\gamma(1)$ gives rise to a well-defined germ of a map from τ at $\gamma(0)$ to τ at $\gamma(1)$ which depends on the path γ only to the extent that it determines the sequence of charts. Thus this germ is unchanged if we vary γ by a small leafwise homotopy keeping endpoints fixed.

A different choice of charts will give rise to the same germ of a map from τ to itself, since transition functions are always cocycles. It follows that there is a well-defined homomorphism from the *fundamental groupoid* of homotopy classes of leafwise paths with endpoints on τ to the groupoid of germs of self-homeomorphisms of τ.

If we restrict to paths which start and end at a fixed $p \in \lambda$ on a transversal τ, we get a well-defined homomorphism from the *fundamental group* $\pi_1(\lambda, p)$ to the *group* of germs of self-homeomorphisms of τ fixing p.

Either of these homomorphisms is known as the *holonomy* of the foliation.

1.3. **Foliated bundles.** Closely related to the idea of holonomy are a class of foliations that arise directly from representations via the *Borel construction*, which we now describe.

Let M be a p-manifold, and F a q-manifold, and suppose we have a representation $\rho : \pi_1(M) \to \mathrm{Homeo}(F)$. We can build a *foliated bundle* E from ρ whose base space is M, and whose fiber is F, as follows.

Let $\tilde{M}$ denote the universal cover of M. The group $\pi_1(M)$ acts on the product $\tilde{M} \times F$ by

$$\alpha(x, f) = (\alpha(x), \rho(\alpha)(f))$$

where $x \to \alpha(x)$ is the deck group action.

Define $E := \tilde{M} \times F/\pi_1(M)$. Projection to the first factor ("vertical" projection) defines a map $E \to M$ whose fibers are copies of F. In other words, E is an F-bundle over M.

The product $\tilde{M} \times F$ has a "horizontal" foliation $\tilde{\mathcal{F}}$ whose leaves are the products. $\tilde{M} \times$ point. The action of $\pi_1(M)$ on $\tilde{M} \times F$ permutes the leaves of $\tilde{\mathcal{F}}$, so it descends to a foliation $\mathcal{F}$ on the bundle E whose leaves are transverse to the fibers.

If μ is a leaf of $\mathcal{F}$, the vertical projection restricts to a covering map $\mu \to M$. If $\tilde{\mu} = \tilde{M} \times f$ is a leaf of $\tilde{\mathcal{F}}$ covering μ for some $f \in F$, then $\tilde{\mu} \to \mu$ is a (universal) covering map, and $\pi_1(\mu)$ is isomorphic to the stabilizer of the leaf $\tilde{M} \times f$, which is the stabilizer of $f \in F$ in the representation ρ.

Example 1.7 (Holonomy). Suppose $\rho : \pi_1(M) \to \mathrm{Homeo}(\mathbb{R}^n)$ fixes 0, and let E be the foliated $\mathbb{R}^n$ bundle over M associated to the Borel construction, with foliation $\mathcal{F}$. Associated to 0 there is a leaf λ of $\mathcal{F}$ which maps homeomorphically to M under the vertical projection. Holonomy defines a homomorphism from $\pi_1(\lambda)$ to the group of germs of homeomorphisms of $\mathbb{R}^n$ at 0. This representation is precisely the germ of ρ at 0 (after we identify $\pi_1(\lambda)$ with $\pi_1(M)$ by vertical projection).

Example 1.8 (Suspension). Let M be a manifold and let $\varphi : M \to M$ be a homeomorphism. The product $M \times [0, 1]$ has a foliation by intervals $m \times [0, 1]$, and these glue together to make up the leaves of the *suspension foliation* on the mapping torus $M_\varphi := M \times [0, 1]/(m, 1) \sim (\varphi(m), 0)$.

We can think of M_φ as the foliated M bundle over S^1 associated to the representation $\pi_1(S^1) \to \mathrm{Homeo}(M)$ which takes the generator to φ.

1.4. **Foliated bundles and smoothness.** It is easy to use the Borel construction to give examples of foliations with nontrivial restrictions on their analytic quality. We give two examples which are important in the theory of 3-manifolds.

1.4.1. *Kopell's Lemma.* Kopell's Lemma [1] is the following:

Theorem 1.9 (Kopell's Lemma)**.** *Let f, g be two commuting C^2 diffeomorphisms of $[0, 1)$ fixing 0.*

If 0 is an isolated fixed point of f, then either it is an isolated fixed point of g or $g = \mathrm{id}$.

I learned the following proof from Navas [2], Thm. 4.1.1.

Proof. Without loss of generality, assume $f(x) < x$ for all $x \in (0, 1)$.

Suppose the theorem is false so that g fixes $p \in (0, 1)$ and define $p_n := f^n(p)$. By assumption g is nontrivial on each $[p_n, p_{n-1}]$. Since 0 is a non-isolated fixed point of g we must have $g'(0) = 1$.

Let $p_1 < u < v < p$. Then by the chain rule and the triangle inequality

$$|\log(f^n)'(v) - \log(f^n)'(u)| \le \sum_{i=1}^{n} |\log f'(f^{i-1}(v)) - \log f'(f^{i-1}(u))|$$

By hypothesis there is an ordering

$$f^{n-1}(u) < f^{n-1}(v) < f^{n-2}(u) < f^{n-2}(v) < \cdots < u < v$$

so we estimate

$$|\log(f^n)'(v) - \log(f^n)'(u)| \le \int_{f^{n-1}(u)}^{v} \left|\frac{f''(s)}{f'(s)}\right| ds \le \int_0^p \left|\frac{f''(s)}{f'(s)}\right| ds = C$$

for some constant C independent of u, v, n. Here is where we use the hypothesis that f is C^2.

Now for any $x \in [p_1, p]$, using the identity $g(x) = f^{-n} g f^n(x)$ and the chain rule, we obtain

$$g'(x) = \frac{(f^n)'(x)}{(f^n)'(f^{-n} g f^n(x))} g'(f^n(x)) = \frac{(f^n)'(x)}{(f^n)'(g(x))} g'(f^n(x))$$

Taking $g'(f^n(x)) \to 1$ as $n \to \infty$ we obtain $|g'(x)| \le e^C$ from the previous estimate.

But now replacing g by g^m we get $|(g^m)'(x)| \le e^C$ with the same constant C *independent* of m. Now, if g is nontrivial on $[p_1, p]$ the powers of g have unbounded derivatives on this interval, giving a contradiction and proving the theorem. □

Assuming this theorem, we can give an example of a foliation which is C^1 but not C^2.

Example 1.10 (Torus leaf). Let f, g be commuting homeomorphisms of $[0, 1]$ fixing 0. Let 0 be an isolated fixed point of f, and suppose $f(x) < x$ on $(0, 1)$. Now suppose $\text{fix}(g) \cap (0, 1)$ is the union of a countable discrete set of points p_n where $f(p_n) = p_{n+1}$. It is possible to conjugate the action of f and g to make them C^1 on $[0, 1]$, but by Kopell's Lemma they cannot be made C^2 near 0.

Associated to this action there is a foliated I bundle $\mathcal{F}$ over a torus T^2 associated to the representation $\pi_1(T) \to \text{Homeo}(I)$ given by identifying the generators of $\pi_1(T) = \mathbb{Z}^2$ with f and g. The total space of the bundle is a product $T^2 \times I$. There is a torus leaf corresponding to the global fixed point 0, and a cylinder leaf corresponding to the p_n, and this cylinder leaf spirals around the torus. All other leaves are planes which spiral around this cylinder.

The foliation in this example has *finite depth* (in fact, depth 2). It has a pair of compact leaves (the tori on the boundary) which are depth 0, then a noncompact leaf which is proper in the complement of the depth 0 leaves (the cylinder) which is depth 1, then a product family of noncompact leaves which are proper in the complement of the depth 0 and 1 leaves (the planes) which are depth 2.

1.4.2. *Thurston Stability Theorem.* Thurston's Stability Theorem [4] is the following:

Theorem 1.11 (Thurston Stability Theorem). *Let G be a finitely generated group of germs of C^1 diffeomorphisms of $\mathbb{R}^n$ fixing 0, and suppose that the derivative homomorphism $G \to \text{GL}(n, \mathbb{R})$ at 0 is trivial.*

Then for any sequence $p_i \to 0$ where the action is nontrivial there are a sequence of linear rescalings of the action near p_i which converge on compact subsets on some subsequence to a nontrivial action of G on $\mathbb{R}^n$ by translations.

Proof. If g is a germ with trivial linear part, then we write $g(x) = x + \tilde{g}(x)$ where $\tilde{g}(x) = o(x)$ and $\tilde{g}'(x) = o(1)$.

Let g, h be two such germs, and let's restrict attention to an open set U where $|\tilde{g}'|, |\tilde{h}'| < \epsilon$. Let $p \in U$ be a point where $\max(|\tilde{g}(p)|, |\tilde{h}(p)|) = \delta > 0$ and such that the ball of radius δ about p is in U.

Then $hg(p) = p + \tilde{g}(p) + \tilde{h}(p + \tilde{g}(p))$. Since $|\tilde{h}'| < \epsilon$ on the straight line from p to $p + \tilde{g}(p)$ we have

$$|\tilde{h}(p + \tilde{g}(p)) - \tilde{h}(p)| < \epsilon|\tilde{g}(p)| \leq \epsilon\delta$$

So $|\tilde{hg}(p) - (\tilde{h}(p) + \tilde{g}(p))| < \epsilon\delta$ which is small compared to $\max(|\tilde{g}(p)|, |\tilde{h}(p)|)$.

So let $g_1, \cdots, g_n$ be generators for G. If we let $\delta_i = \max_j(|\tilde{g_j}(p_i)|)$ then there is a subsequence for which the vector $\{\tilde{g_j}(p_i)/\delta_i\}$ converges to a nontrivial vector $\{v_j\}$ with $\max_j(|v_j|) = 1$.

Fix a k. Then for any ϵ, let U be the open set where $|\tilde{g}_j'| < \epsilon$ for all j. If i is such that the ball of radius $k\delta_i$ about p_i is contained in U, then if

$$w := g_{i_1}^{\epsilon_1} g_{i_2}^{\epsilon_2} \cdots g_{i_k}^{\epsilon_k}$$

is a word in the generators of length $\leq k$ (here each $\epsilon_l = \pm 1$) we have

$$|\tilde{w} - \sum_l \epsilon_l \tilde{g_{i_l}}| < \epsilon\delta_i k$$

which is small compared to δ_i, at least for fixed k, and for ϵ small depending on k. It follows that linear rescalings of the action of G at the p_i by $1/\delta_i$ converges on finite subsets of G and compact subsets of $\mathbb{R}^n$ to an action by translations, where g_j converges to the translation by v_j. □

As a corollary we deduce that a group G acting as in the hypothesis of the theorem must have $\mathrm{Hom}(G; \mathbb{R}) \neq 0$.

Example 1.12. Let λ be a once-punctured torus. Then $\pi_1(\lambda) = F_2$ generated by elements a, b. Pick a complete hyperbolic structure on λ, and let $\rho : \pi_1(\lambda) \to \mathrm{Homeo}(S^1)$ be the induced action on the circle at infinity.

Since $\pi_1(\lambda)$ is free, we can lift ρ to a representation $\tilde{\rho} : \pi_1(\lambda) \to \mathrm{Homeo}(\mathbb{R})$ in which both a and b fix infinitely many points accumulating at the ends. Notice that the commutator $[a, b]$ acts freely on $\mathbb{R}$.

Now identify $\mathbb{R}$ with $(0, 1)$ and extend $\tilde{\rho}$ trivially to the endpoints to get a representation $\sigma : \pi_1(\lambda) \to \mathrm{Homeo}(I)$, and let E be the associated foliated I bundle over λ, and $\mathcal{F}$ the foliation

Suppose σ could be conjugated to be C^1. Since the generators a, b have fixed points accumulating to 0 their derivatives must be equal to 1 there, so the linearization of the action at 0 is trivial. But then the Thurston Stability Theorem would imply that linear rescalings of the action at a sequence of points $p_i \to 0$ would converge to an action by translations. In particular, one of a or b would move points near 0 "more" than the commutator $[a, b]$,

contrary to the definition of σ. This contradiction shows that σ cannot be made C^1, and therefore neither can $\mathcal{F}$.

1.5. **Codimension one.** There are a number of special features of codimension 1 foliations that will be very important in what follows.

Proposition 1.13 (Nonclosed leaf, closed transversal)**.** *Let $\mathcal{F}$ be a codimension 1 foliation of a closed manifold M, and let λ be a leaf which is not closed. Then there is a closed transversal γ which intersects λ.*

Proof. Notice that in a closed manifold a leaf is nonclosed if and only if it is noncompact. But then there is some product chart to which the leaf recurs infinitely often, and consequently it accumulates somewhere.

It follows that there is some transverse interval I which starts and ends on λ and has a consistent co-orientation at both endpoints. Join the endpoints of I by a path J in λ. The union is a closed loop. By the co-orientation condition, we can perturb this loop to be transverse to $\mathcal{F}$ and still to intersect λ. See Figure 1.

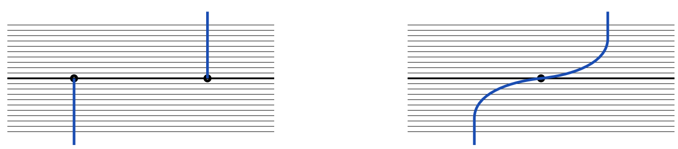

FIGURE 1. A nonclosed leaf admits a transversal

□

The next Theorem is due to Novikov, and is a cornerstone of the codimension 1 theory.

Theorem 1.14 (Novikov, Closed leaves are closed)**.** *Let $\mathcal{F}$ be a codimension 1 foliation of a closed manifold M. Then the union of the closed leaves of $\mathcal{F}$ is closed.*

Proof. After passing to a finite cover we can assume $\mathcal{F}$ is oriented and co-oriented. Taking finite covers preserves the property of being closed, so it suffices to prove the theorem in the cover.

Since M is closed, $H_{n-1}(M)$ is finite-dimensional, and therefore so is the subspace of $H_{n-1}(M)$ generated by closed leaves. So there are finitely many closed leaves $\lambda_1, \cdots, \lambda_m$ so that any other closed leaf is homologous to a linear combination of the λ_i.

Let μ_i be a sequence of closed leaves which have μ in their limit, and suppose μ is not closed. Then certainly μ is disjoint from the λ_i, and by Proposition 1.13 we can find a closed transversal γ which intersects μ and does not intersect any λ_i (just look at a chart in which μ accumulates and construct I as in the proposition disjoint from the intersection of the chart with the λ_i). But then γ intersects μ_j transversely for big enough j, and since $\mathcal{F}$ is co-oriented, this intersection is homologically essential. Thus, $[\mu_j]$ is nontrivial in $H_{n-1}(M)$, but is not in the span of the $[\lambda_i]$, which is a contradiction. □

Codimension 1 is essential for this theorem, as the following example shows.

Example 1.15. Let $\varphi : D \to D$ take the unit disk to itself by rotating the circle at radius r by $2\pi r$. The suspension of φ gives a codimension 2 foliation of a closed solid torus. The closed leaves are the suspensions of the circles at rational radius. Thus the closed leaves are dense but not closed.

Lemma 1.16. *Let $\mathcal{F}$ be a codimension 1 foliation of a closed manifold M. Suppose that $\mathcal{F}$ is co-oriented. Let λ be a closed leaf which is a limit of closed leaves λ_i. Then the λ_i are homeomorphic to λ.*

Proof. Foliate a product neighborhood U of λ by transversals, and choose a basepoint $p \in \lambda$ on a transversal τ. Because λ is closed, there is a compact family of paths Γ in λ starting at p and ending at every other point in λ. For example, we can choose a Riemannian metric on λ, and let Γ consist of the paths with length equal to at most half the injectivity radius. Consequently there is a transversal σ contained in τ so that holonomy transport along any path in Γ is defined on all of σ, and takes it into a transversal in the given product.

For large enough j we must have $\lambda_j \cap \sigma$ nonempty. Let q be such a point. For each $r \in \lambda$ intersecting the transversal $\tau(r)$ there is a path $\gamma(r)$ in Γ from p to r and holonomy transport takes q to $q(r) \in \tau(r)$. For points r with more than two paths $\gamma(r), \gamma'(r)$ in Γ (e.g. for r on the cut locus of p for the choice of Γ given above) we might a priori have two $q(r), q'(r) \in \tau(r)$. In fact we claim $q(r) = q'(r)$.

For, otherwise, $q'(r) < q(r)$ (say), so holonomy around the loop $\beta := \gamma'(r) \circ \gamma(r)^{-1}$ contracts the interval $[r, q(r)]$ to $[r, q'(r)]$ by a homeomorphism h. But then $q(r) > h(q(r)) > h^2(q(r)) > h^3(q(r)) > \cdots$ and so on, so that λ_j contains infinitely many points on $\tau(r)$, contrary to the fact that λ_j is closed. This proves the claim.

Thus the function $r \to q(r)$ defines a bijection from λ to λ_j, whose inverse is given by projection along the transversals, which gives the desired homeomorphism from λ_j to λ. □

If $\mathcal{F}$ is not co-oriented, then a one-sided closed leaf λ might be a limit of closed leaves λ_i that double cover it.

Theorem 1.17 (Reeb stability)**.** *Let $\mathcal{F}$ be a codimension 1 foliation of a closed connected manifold M. Suppose some closed leaf λ has π_1 finite. Then M is finitely covered by a $\tilde{\lambda}$ bundle over S^1 foliated by fibers.*

Proof. Pass to a finite cover where $\pi_1(\lambda) = 1$ and $\mathcal{F}$ is co-oriented. Since λ is simply-connected, holonomy transport is trivial where defined. Since it is closed, we can foliate some neighborhood as a product. Thus the set of leaves homeomorphic to λ is nonempty and open.

But by Theorem 1.14 and Lemma 1.16 the set of leaves homeomorphic to λ is also closed, so it is all of M, and we see that the structure is locally that of a product, and globally that of a bundle. □

In particular, if M is a closed 3-manifold, and $\mathcal{F}$ is a 2-dimensional foliation with an S^2 leaf, then M is finitely covered by $S^2 \times S^1$ and $\mathcal{F}$ is covered by the product foliation by spheres.

2. Reeb components and Novikov's Theorem

From now on we focus exclusively on cooriented codimension one foliations of oriented 3-manifolds, unless we explicitly say to the contrary. Furthermore, we assume that our foliations have no spherical leaves, since the Reeb Stability Theorem 1.17 says that the only (coorientable) foliation with a spherical leaf is the product foliation of $S^2 \times S^1$.

Having ruled out spherical leaves, we next consider toral ones.

2.1. **Reeb components.** Let $\mathbb{R}^3_+$ denote the closed upper half-space in $\mathbb{R}^3$; i.e. the subset where $z \geq 0$. Let $\mathbb{R}^3_+ - 0$ denote the complement of the origin in $\mathbb{R}^3_+$. Let $\tilde{\mathcal{F}}$ be the foliation of $\mathbb{R}^3_+ - 0$ by horizontal leaves $z =$ constant. Note that the leaves with $z > 0$ are all planes, but $z = 0$ is a punctured plane.

The diffeomorphism $\varphi : (x, y, z) \to (2x, 2y, 2z)$ takes $\mathbb{R}^3_+ - 0$ to itself and permutes the leaves of $\tilde{\mathcal{F}}$. Therefore it descends to a foliation $\mathcal{F}$ of the quotient $(\mathbb{R}^3_+ - 0)/\langle\varphi\rangle$ which is homeomorphic to a solid torus $D^2 \times S^1$.

The foliation $\mathcal{F}$ has one torus leaf which is the boundary $S^1 \times S^1$, and is covered by the leaf $z = 0$ of $\tilde{\mathcal{F}}$. All other leaves are planes, whose end winds around the torus leaf, like infinitely deep "socks" with their toes stuffed into their mouths.

Definition 2.1. The foliation $\mathcal{F}$ of the solid torus is called the *Reeb foliation*. A solid torus in a foliated 3-manifold with such a foliation is called a *Reeb component*.

A foliation of a 3-manifold is *Reebless* if it has no Reeb component.

Example 2.2. Any Lens space (for instance, S^3) admits a foliation obtained by gluing two Reeb components along their boundaries.

Note that we can define "Reeb components" of $S^1 \times D^{n-1}$ for any n by replacing 3 by n in the construction above. To distinguish these we call these *n-dimensional Reeb components*.

2.2. **Turbularization.** Reeb components can be "inserted" into foliations along transverse loops. Suppose $\mathcal{F}$ is a foliation of M and γ is an embedded loop transverse to $\mathcal{F}$. Remove a solid torus neighborhood $N(\gamma)$ from $\mathcal{F}$. Then $M - N(\gamma)$ has boundary a torus T, and we look at a product collar $T \times [0, 1]$, where $T \times 0 = \partial(M - N(\gamma))$ and if we write $T = S^1 \times S^1$ then the leaves of $\mathcal{F}$ are transverse to the $S^1 \times$ point factors (we call this first S^1 factor "vertical"; it is the direction of the loop γ).

For each n we let φ_n be a diffeomorphism of $T \times [0, 2^{-n}]$ which is the identity on the boundary and which takes the horizontal point $\times$ point $\times$ $[0, 2^{-n}]$ factors and drags them around the vertical S^1 factor. This "spins" the leaves of $\mathcal{F}$ once around the vertical direction. The composition $\varphi := \prod_{i=0}^{\infty} \varphi_n$ is smooth in the interior, and the leaves of $\varphi(\mathcal{F})$ accumulate on the boundary T^2. We can therefore define a new foliation $\mathcal{F}'$ of M which agrees with $\varphi(\mathcal{F})$ on $M - N(\gamma)$, and which has a Reeb component on $N(\gamma)$.

We say that $\varphi(\mathcal{F})$ on $M - N(\gamma)$ is obtained from $\mathcal{F}$ on $M - N(\gamma)$ by *spinning* leaves along the boundary, and call $\mathcal{F}'$ the result of *turbularization* of $\mathcal{F}$ along γ.

Note that $T\mathcal{F}$ and $T\mathcal{F}'$ are homotopic as plane fields.

2.3. **Constructing foliations.** So far we have not given many examples of 3-manifolds with foliations. The following construction is due to Thurston.

Theorem 2.3 (Constructing foliations)**.** *For any 3-manifold, every homotopy class of 2-plane field is homotopic to $T\mathcal{F}$ for some foliation $\mathcal{F}$.*

Proof. On a small enough scale the 2-plane field is "almost" constant, and looks like the horizontal distribution on $\mathbb{R}^3$. Choose a fine triangulation whose simplices are very close to linear on this small scale, and whose 1-skeleton is transverse to the 2-plane field. Locally, where the 2-plane field can be co-oriented, each simplex inherits a total order on its vertices, and its edges can be oriented so that each edge points to the higher of its two terminal vertices.

After barycentric subdivision if necessary, the simplices can be 2-colored black and white so that adjacent simplices have different colors. Each simplex has a highest and a lowest vertex; the boundary is a sphere, and we give each sphere a (singular) foliation which spirals from the lowest to the highest vertex, where the spiraling is clockwise on the white simplices and anticlockwise on the black simplices. Tilting the foliation in the anticlockwise direction on a face of a black simplex tilts it in the clockwise direction as seen from an adjacent white simplex; these deformations therefore interfere "constructively", and the desired foliation can be achieved. Then extend the resulting foliation to an open neighborhood N of the 2-skeleton.

The interior of each simplex is a ball; the foliation on the boundary of each simplex typically cannot be extended to the interior, because of the spiraling.

However we claim that there is a positively oriented transversal in N from any point to any other point; to see this, observe that in the boundary of a simplex σ where the foliation spirals from bottom to top, a transverse path that spirals in the opposite direction can move (almost) from top to bottom. By moving from simplex to adjacent simplex, we can move anywhere in the manifold. This proves the claim.

For each simplex σ by this claim we can drill out a positively oriented transverse path τ from the top of σ to the bottom. Then a neighborhood of τ together with the interior of σ is a solid torus. We can spin the foliation around the boundary torus, and then fill it in with a Reeb component. Doing this for each simplex produces a foliation in the desired homotopy class. □

2.4. **Novikov's Theorem.**

Theorem 2.4 (Novikov Reebless)**.** *Let M be a 3-manifold and let $\mathcal{F}$ be a Reebless foliation. Then*

(1) every leaf λ is π_1-injective; and
(2) every transverse loop γ is essential in π_1.

Proof. Suppose not, so that there is some homotopically trivial loop γ which is either transverse or tangent to $\mathcal{F}$. Let D be a disk that γ bounds. Put D in general position relative to $\mathcal{F}$.

Thus D inherits a singular foliation. Because ∂D is either transverse or tangent to $\mathcal{F}$, and because $\chi(D) = 1$ it follows that there is a point $p \in D$ which looks like a local minimum/maximum singularity with respect to $\mathcal{F}$. Thus there is a maximal open neighborhood U of p foliated by concentric circles S_t for $t \in (0, 1)$ where S_t bounds a disk E_t in its leaf λ_t for $t \in (0, 1)$. Let S_1 be the limit of the S_t. By hypothesis, S_1 is a (possibly singular) circle in some leaf λ_1.

If S_1 bounds a disk E_1 which is a limit of the E_t then the holonomy around S_1 is trivial, so either we can extend U (contrary to maximality) or there is a singularity on S_1. In the latter case we can push U into E_1 and cancel a pair of singularities, reducing the complexity.

If not, then the E_t do not converge on compact subsets. This is only possible if their areas increase without bound. We must show in this case that there is a Reeb component.

First observe that by cut-and-paste we can restrict to the case that each E_t is embedded; for, leafwise we can reduce the number of self-intersections of S_t by a homotopy, and these homotopies can be performed in a family unless some innermost family of embedded disks has unbounded area.

For small t the union $B_t := \cup_{s\in[0,t]} E_s$ is a closed ball bounded by $E_t \cup U|_{[0,t]}$. As we increase t this ball expands. Since the areas of the E_t increase without limit, the volume swept out by this family must increase without limit also, and since M is compact, eventually B_t must intersect itself. This can only be because some E_t becomes tangent to $U|_{[0,t]}$, and this can only be at the center point p, because this is the only place where U is tangent to $\mathcal{F}$. Thereafter B_t is an expanding family of solid tori which can never develop any more self-tangencies. But this means that the volume and diameter of this solid torus are a priori bounded (by that of M) and they must limit to a solid torus which by construction is foliated as a Reeb component.

So if $\mathcal{F}$ is Reebless, we may inductively cancel all the singularities and push D into a leaf of λ (showing that a homotopically trivial tangential loop γ is inessential in its leaf) or arrive at a contradiction (showing that a transverse loop γ is homotopically essential after all). □

Corollary 2.5. *Let M be a 3-manifold and let $\mathcal{F}$ be Reebless without spherical leaves. Then if $\tilde{\mathcal{F}}$ denotes the foliation of the universal cover $\tilde{M}$, leaves of $\tilde{\mathcal{F}}$ are properly embedded planes, and M is irreducible.*

Proof. If some leaf λ of $\tilde{\mathcal{F}}$ is not proper, it accumulates somewhere, and we can build a transverse loop γ to $\tilde{\mathcal{F}}$ which projects to an inessential transverse loop in M, contrary to Theorem 2.4.

Again, by Theorem 2.4, leaves of $\tilde{\mathcal{F}}$ are simply-connected, so (since $\mathcal{F}$ has no spherical leaves) they are all planes.

Alexander's proof of the irreducibility of $\mathbb{R}^3$ depends only on the fact that it has a foliation by (proper) planes. The same argument shows that $\tilde{M}$ is irreducible, and therefore so is M. □

3. Taut foliations

In this section we introduce the class of *taut foliations*. These are the foliations which see and certify the most interesting geometric and topological properties of their ambient manifold, and are the focus of the remainder of these notes.

For the sake of simplicity we're going to assume (for now) that our foliations are smooth. In fact, this is a substantial simplification, unwarranted in many important situations. Fortunately, there is a version of the theory that makes sense with no assumptions of regularity, and for which all the most important theorems and appliations still go through; we defer the discussion of this to § 4.

3.1. **Equivalent formulations of tautness.** If $\mathcal{F}$ is a codimension 1 foliation of a 3-manifold M, after passing to a cover of degree at most 4 we may assume that M is oriented, and that $\mathcal{F}$ is oriented and co-oriented. This means that there are orientations on the tangent bundle $T\mathcal{F}$ and the normal bundle $\nu\mathcal{F}$ of $\mathcal{F}$ respectively which together give an orientation on $T\mathcal{F} \oplus \nu\mathcal{F} = TM$ agreeing with the given orientation on M.

Theorem 3.1 (Equivalent Formulations of Tautness)**.** *For a smooth, oriented, co-oriented codimension 1 foliation $\mathcal{F}$ of a connected closed 3-manifold M the following conditions are equivalent:*

1. *(1) For every point $p \in M$ there is an immersed circle $\gamma_p : S^1 \to M$ transverse to $\mathcal{F}$ and passing through p.*
2. *(2) There is an immersed circle $\gamma : S^1 \to M$ transverse to $\mathcal{F}$ and intersecting every leaf of $\mathcal{F}$.*
3. *(3) There is no proper compact submanifold N of M whose boundary is tangent to $\mathcal{F}$, and for which the co-orientation points in to N along ∂N.*
4. *(4) There is a closed 2-form ω on M^3 positive on $T\mathcal{F}$.*
5. *(5) There is a flow X transverse to $\mathcal{F}$ which is volume preserving for some Riemannian metric on M.*
6. *(6) There is a Riemannian metric on M for which every leaf of $\mathcal{F}$ is a minimal surface.*
7. *(7) There is a Riemannian metric on M and a closed 2-form ω which calibrates $\mathcal{F}$.*

An $\mathcal{F}$ satisfying any of these equivalent conditions is said to be taut*.*

If $\mathcal{F}$ is orientable and coorientable and contains a spherical leaf (i.e. a leaf homeomorphic to $\mathbb{RP}^2$ or S^2), then the Reeb Stability Theorem 1.17 says that $\mathcal{F}$ is the product foliation of $S^2 \times S^1$ by spheres. This foliation evidently satisfies all the proposed definitions of tautness, so for the remainder of this section we'll assume without explicitly saying so that $\mathcal{F}$ has no spherical leaves.

We prove Theorem 3.1 in a series of steps.

3.1.1. *Transversals.* An immersed circle $\gamma : S^1 \to M$ transverse to $\mathcal{F}$ is called a *transverse loop*, or a *transversal*. Transversals can be made embedded by a small leafwise homotopy.

Corollary 3.2. *Being taut is inherited under passing to finite covers.*

Proof. A connected preimage of a closed transversal is a closed transversal. □

If a transversal passes through a given point p on a leaf λ, it's useful to be able to modify it by leafwise homotopy so that it passes through another given point q on λ. This can be arranged:

Lemma 3.3 (Move transversal)**.** *Suppose $\gamma : S^1 \to M$ is transverse to $\mathcal{F}$. Suppose $\gamma(0) = p$ contained in a leaf λ of $\mathcal{F}$, and let q be any other point on λ. Then we may homotop γ, through maps transverse to $\mathcal{F}$, to a new map with $\gamma(0) = q$.*

Proof. Let α be an embedded path in λ from p to q. A sufficiently small neighborhood U of α is foliated as a product in such a way that the image of γ intersects U in a single vertical segment. We may interpret this segment as the graph of a (constant) map from $[-1, 1]$ to λ starting and ending at p. This map is homotopic to α concatenated with α^{-1}; the graph of this homotopy gives the desired modification of γ. See Figure 2.

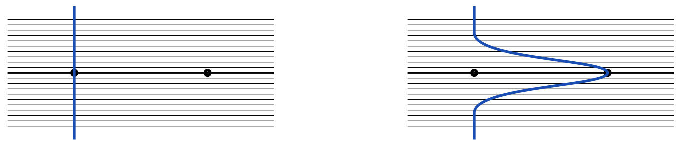

FIGURE 2. Leafwise homotop γ so it passes through q

□

Corollary 3.4. *A foliation without closed leaves is taut.*

Proof. By Proposition 1.13 every nonclosed leaf admits a transversal, and by Lemma 3.3 we can find a transversal passing through any given point on such a leaf. □

Claim. *(1) and (2) are equivalent.*

Proof. Suppose $\gamma : S^1 \to M$ is transverse to $\mathcal{F}$. Let $U(\gamma)$ be the union of the leaves of $\mathcal{F}$ that γ intersects. Evidently $U(\gamma)$ is open. Furthermore, by Lemma 3.3, for any $p \in U(\gamma)$ we may modify γ by a leafwise homotopy to some new δ so that $U(\gamma) = U(\delta)$ and δ passes through p. This shows that (2) implies (1).

To see that (1) implies (2), hypothesis (1) says that for every $p \in M$ there's γ with p in (the image of) γ, and this means $p \in U(\gamma)$. Since M is compact, there's a smallest (finite) collection of γ_i for which the union of the $U(\gamma_i)$ is equal to M.

We claim that the $U(\gamma_i)$ are disjoint. Otherwise there would be i and j with $p \in U(\gamma_i) \cap U(\gamma_j)$. By Lemma 3.3 we could modify γ_i and γ_j by a leafwise homotopy so they both pass through p. Then we could build a new transversal γ by starting at p, first going around γ_i, and then going around γ_j. This would satisfy $U(\gamma) = U(\gamma_i) \cup U(\gamma_j)$, which contradicts our choice of transversals to be as few as possible.

So: the $U(\gamma_i)$ are disjoint and open. Since M is connected, there can be only one transversal in the collection. □

3.1.2. *Dead ends.* A compact submanifiold N of M whose boundary is tangent to $\mathcal{F}$ and for which the co-orientation points inwards along ∂N is called a *dead end*. Formulation (3) says that tautness is equivalent to having no dead ends.

Lemma 3.5 (Boundary tori)**.** *Let N be a dead end component. Then ∂N is a union of tori.*

Proof. We show $\chi(\partial N) = 0$. Since ∂N has no spheres (by fiat) it follows that all the components will be tori.

The coorientation on $\mathcal{F}$ lets us find a nonsingular vector field X everywhere transverse to $\mathcal{F}$ and pointing inwards along ∂N. Let Y be a generic vector field tangent to ∂N and extend it as a product on a collar of ∂N. Then let Z be equal to X away from this collar, and on the collar equal to a convex combination of X and Y, limiting to Y on ∂N. We double Z to get a vector field on DN singular only at the singular points of Y, and whose singularities have the same index on DN as on ∂N. Thus $\chi(DN) = \chi(\partial N) = 0$. □

Claim. *(1) and (3) are equivalent.*

Proof. A transversal which enters a dead end component can never leave, so any foliation with a dead end does not satisfy (1).

Conversely, suppose $\mathcal{F}$ does not satisfy (1), and let λ be a (necessarily compact) leaf which does not intersect a closed transversal.

Let N be the subset of M consisting of points which can be reached by a positively oriented transversal starting at λ. Transversals can always be extended, so N is open and a union of leaves. Thus its closure $\bar{N}$ is compact and has boundary a union of leaves. Moreover the coorientation points inwards all along $\bar{N}$, since otherwise a positively oriented transversal from a point near $\partial\bar{N}$ could be extended all the way to $\partial\bar{N}$.

By construction $\bar{N}$ is a dead end (with λ as a leaf). □

Taut foliations are therefore Reebless, but the converse is false.

Example 3.6. A pair of (oppositely aligned) 2-dimensional Reeb components foliates a torus (with 1-dimensional leaves). Taking the product with S^1 produces a foliation of a 3-torus which is Reebless but not taut. There are two dead end components, each equal to the product of a 2-dimensional Reeb component with a circle, and homeomorphic to a torus times interval.

Corollary 3.7. *If M is hyperbolic, $\mathcal{F}$ is Reebless if and only if it is taut.*

Proof. If $\mathcal{F}$ is Reebless, every leaf is essential, by Novikov's Theorem 2.4. So if M admits a foliation which is Reebless but not taut, then M is toroidal. □

3.1.3. *Forms and flows.* We have proved the equivalence of formulations (1)–(3). We now prove the equivalence of formulations (4)–(7) with themselves and with (1)–(3).

Since we're assuming throughout this section that our foliations $\mathcal{F}$ are smooth and co-oriented, we can find some nowhere zero 1-form α so that $T\mathcal{F} = \ker\alpha$.

Claim. *(1) implies (4).*

Proof. For every point $p \in M$ there's a transversal $\gamma : S^1 \to M$ through p. By a homotopy we may assume γ is smooth and embedded. Thus, an open neighborhood of $\gamma(S^1)$ is an open solid torus N whose induced foliation is a product $D^2 \times S^1$. Since M is compact, finitely many N_i cover M.

Let θ be a positive 2-form on D^2 that is nowhere zero, and tapers smoothly to zero at the boundary. Each solid torus N_i projects to the D^2 factor, and we can pull back θ to a closed 2-form ω_i on N_i, and then extend it to zero outside N_i. Let $\omega = \sum_i \omega_i$. Then ω is closed and positive on $T\mathcal{F}$. □

Claim. *(4) implies (5).*

Proof. If ω and α are as above, then $\omega \wedge \alpha$ is a nowhere vanishing 3-form on M; i.e. a volume form. Since ω is nondegenerate, $\ker(\omega)$ is 1-dimensional everywhere, and transverse to $\ker(\alpha) = T\mathcal{F}$. Thus there is a nowhere zero vector field X, transverse to $\mathcal{F}$, satisfying $\iota_X\omega = 0$ and $\alpha(X) = 1$ everywhere. We may give M a Riemannian metric in which X has length 1 and is perpendicular to $\mathcal{F}$, and in which ω restricts to the area form on $T\mathcal{F}$. For such a Riemannian metric, the volume form is $\omega \wedge \alpha$.

But then Cartan's formula gives

$$\mathcal{L}_X\omega \wedge \alpha = d\iota_X(\omega \wedge \alpha) = d\omega = 0$$

so that X generates a *volume-preserving flow* transverse to $\mathcal{F}$. □

Claim. *(5) implies (3).*

Proof. If there is a volume preserving transverse flow, there can be no dead end, since any transverse flow would take a dead end properly inside itself, thereby compressing it. □

A closed 2-form ω on a Riemannian 3-manifold M *calibrates* a surface S if $\omega(\xi) \leq \text{area}(\xi)$ for every 2-plane ξ, and if $\omega(\xi) = \text{area}(\xi)$ for (oriented) 2-planes ξ tangent to $\mathcal{F}$.

Claim. *(4) implies (7).*

Proof. In the Riemannian metric for which $\omega \wedge \alpha$ is a volume form and X is perpendicular to $T\mathcal{F}$, the form ω is calibrating for $\mathcal{F}$. □

A surface S in a 3-manifold M is said to be *least area in its homology class* if for any compact subsurface $D \subset S$ and any other $D' \subset M$ with $\partial D' = \partial D$ and D' homologous to D rel. boundary, we have

$$\text{area}(D') \geq \text{area}(D)$$

If S is compact we allow $D = S$ here.

A surface which is least area in its homology class is certainly least area in its isotopy class, and in particular it is a stable minimal surface.

Lemma 3.8 (Calibrated is least area)**.** *Suppose ω calibrates S. Then S is least area in its homology class.*

Proof. Let $D \subset S$ be compact and $D' \subset M$ have $\partial D = \partial D'$ and D' homologous to D rel. boundary. Then

$$\text{area}(D') \geq \int_{D'} \omega = \int_D \omega = \text{area}(D)$$

□

Claim. *(7) implies (6).*

Proof. This is an immediate corollary of Lemma 3.8. □

Claim. *(6) implies (5).*

Proof. Let M be a Riemannian manifold for which the leaves of $\mathcal{F}$ are minimal surfaces. Let X be an orthogonal vector field of constant length 1. Let φ_t be the flow generated by X. Since X is perpendicular to $T\mathcal{F}$, the tangent field $\varphi_{-t}^* T\mathcal{F}$ stays perpendicular to X to first order in t. It follows that the first variation of the volume is equal to the first variation of the area of leaves under the flow. But this first variation of area under orthogonal flow is (by definition) the mean curvature, and a surface is minimal if and only if its mean curvature vanishes identically. □

This completes the proof of Theorem 3.1.

3.2. **Invariant transverse measures.** Let $\mathcal{F}$ be a foliation. A *invariant transverse measure* μ is a measure on the local leaf space in each chart which is preserved by transition functions.

Said another way, it assigns a non-negative number $\mu(\tau)$ to each transversal, which is (countably) additive, and so that if τ' can be obtained from τ by a leafwise homotopy, then $\mu(\tau') = \mu(\tau)$.

Let μ be a (nontrivial) transverse measure. Let τ be a total transversal; i.e. τ consists of a finite union of intervals $\tau_1 \cup \cdots \cup \tau_k$. Then μ gives rise to a measure on τ, and holonomy transport preserves this measure.

If $\mathcal{F}$ is codimension 1 and coorientable, each τ_i is an interval. Integrating μ along τ_i (in the positive direction) gives rise to a monotone map from τ_i to an interval σ_i with $\text{length}(\sigma_i) = \mu(\tau_i)$. This map might not be continuous (if μ has atoms) but the holonomy action on τ induces an action on $\sigma := \cup_i \sigma_i$ by (germs of) *translations*, and the action of holonomy transport on σ completely encodes the action of holonomy transport on the leaves of $\mathcal{F}$ in the support of μ.

If M is compact, the groupoid is compactly generated by leafwise paths of uniformly bounded length; we could equally well talk about a finitely generated *pseudogroup* of partially defined translations of σ.

Here is another way to think about this pseudogroup. Since $\mathcal{F}$ is coorientable, we can extend μ to a *signed* measure on *oriented* transversals, where changing the orientation of a transversal gives the negative of the measure.

A generic smooth path γ in M can be decomposed into a finite union of positive and negative transversals, and we can define $\mu(\gamma)$ by additivity. A generic smooth homotopy of γ rel. endpoints creates or destroys transversal segments in (leafwise homotopic) pairs with opposite orientations, so μ is invariant under such homotopies. Thus we obtain a *homomorphism*

$$\rho_\mu : \pi_1(M) \to \mathbb{R}$$

Let λ be a leaf in the support of μ. Suppose λ is not compact, so that μ has no atoms on λ. We can consider the returns of λ to a product chart U. Locally the leaf space is parameterized by a transversal τ, and we can fix a point $p \in \lambda \cap \tau$. A leafwise path $\gamma \subset \lambda$ from p to $q \in \tau$ can be closed up with the oriented interval $[q,p] \subset \tau$ to make a closed loop γ', and $\mu([q,p]) = \rho_\mu(\gamma')$. Since λ is in the support of μ, the point q is determined by p and the (signed) value of $\mu([q,p])$. Thus the number of points of $\lambda \cap \tau$ in the ball of radius R in λ about p is bounded by the number of values of ρ_μ on the set of loops in $\pi_1(M)$ with representatives of length at most $R + \text{length}(\tau)$. Since $\mathbb{R}$ is abelian, the latter grows polynomially with degree at most $b_1(M)$.

We deduce the following theorem of Plante [3]:

Theorem 3.9 (Plante polynomial growth)**.** *Let $\mathcal{F}$ be a codimension 1 foliation of M. If λ is a leaf in the support of an invariant transverse measure μ then λ has polynomial growth of degree $\leq b_1(M)$.*

In fact, there is a (partial) converse to this theorem, coming from the amenability of pseudogroups of subexponential growth.

Proposition 3.10 (subexponential growth)**.** *Let $\mathcal{F}$ be a foliation of a compact manifold M, and let λ be a leaf with subexponential growth. Then there is a nontrivial invariant measure on $\mathcal{F}$ with support contained in the closure of λ.*

Proof. Let λ_i be a sequence of Følner subsets of λ; i.e. subsets with

$$\text{volume}(\partial\lambda_i)/\text{volume}(\lambda_i) \to 0$$

For each transversal τ we can define

$$\mu_i(\tau) = \#(\tau \cap \lambda_i)/\text{volume}(\lambda_i)$$

If τ and τ' are obtained by leafwise homotopy of bounded length, the Følner condition guarantees $|\mu_i(\tau') - \mu_i(\tau)| \to 0$. No transversal τ can intersect any λ_i in more than $C(\tau)\text{volume}(\lambda_i)$ points, and since M is compact, a total transversal will intersect λ_i for big i in $\epsilon \cdot \text{volume}(\lambda_i)$ points.

Thus some subsequence of the measures μ_i converges to a nontrivial invariant transverse measure, whose support is contained in the closure of λ. □

Example 3.11. The planar leaves of a Reeb component have linear growth. The construction from Proposition 3.10 gives rise to an (atomic) invariant measure supported on the boundary torus.

When $\mathcal{F}$ is coorientable, we can think of the weighted subsets $\lambda_i/\text{volume}(\lambda_i)$ as (singular) p-chains (if p is the dimension of the leaves), and their limit as a *de Rham p-cycle* representing a p-dimensional homology class $[\mu]$. When $\mathcal{F}$ has codimension one, $[\mu]$ is (Poincaré) dual to a class in $H^1(M;\mathbb{R}) = \text{Hom}(\pi_1(M);\mathbb{R})$; evidently this is the class of the homomorphism ρ_μ.

Using the formalism of transverse measures, we can give another more refined characterization of tautness.

Proposition 3.12 (Taut and transverse measures)**.** *Let M be a compact 3-manifold. Let $\mathcal{F}$ be codimension 1 and cooriented. Then there is a closed 2-form ω which is positive on $T\mathcal{F}$ and in the cohomology class $[\omega]$ if and only if $[\omega]([\mu]) > 0$ for all invariant transverse measures μ.*

Proof. If μ is a transverse measure which is the limit of μ_i associated to subsets λ_i of leaves of $\mathcal{F}$, the pairing $[\omega]([\mu])$ is the limit

$$[\omega]([\mu]) = \lim_{i\to\infty} \frac{1}{\text{volume}(\lambda_i)} \int_{\lambda_i} \omega$$

so we must necessarily have $[\omega]([\mu]) > 0$ for every such measure.

Conversely, suppose there is a cohomology class $[\omega]$ with $[\omega]([\mu]) > 0$ for all invariant transverse measures μ. This gives rise to a linear functional on the space of de Rham cycles which is strictly positive on the cone of cycles represented by invariant transverse measures. By the Hahn-Banach theorem this extends to a linear functional on the space of all de Rham chains and therefore defines a cocycle. By convolution we can approximate this functional by an honest closed form ω in the class of $[\omega]$. □

One can think of this just as easily in terms of volume preserving flows. A volume preserving flow X can be thought of as a de Rham 1-cycle. The intersection pairing with the homology class of a 2-cycle S is equal to the flux of X through S:

$$[X] \cap [S] = \text{flux of } X \text{ through } S$$

An equivalent statement of the proposition is that a 1-dimensional homology class $[X]$ is represented by a volume preserving transverse flow X if and only if $[X] \cap [\mu] > 0$ for all invariant transverse measures μ.

Example 3.13. A dead end component is bounded by a union of compact tori which collectively represent zero in homology. An atomic transverse measure with equal weight on each torus represents zero in homology, and obstructs the existence of a form ω.

4. Smoothness and some substitutes

5. Holomorphic geometry

5.1. **Uniformization.**

Theorem 5.1 (Candel, Uniformization)**.** *M has a metric for which every leaf of $\mathcal{F}$ is hyperbolic if and only if every invariant transverse measure μ has $\chi(\mu) < 0$.*

5.2. **Poincaré Series.**

Theorem 5.2 (Ghys, Ample)**.** *A smooth co-orientable foliation $\mathcal{F}$ of M is taut if and only if there is some integer n, and a map $\varphi : M \to \mathbb{CP}^n$ which restricts to a holomorphic immersion on each leaf.*

Proof. Suppose there is such a φ. The pullback of the Kähler form on $\mathbb{CP}^n$ is a closed 2-form on M, positive on $T\mathcal{F}$, so $\mathcal{F}$ is taut.

Conversely, suppose $\mathcal{F}$ is taut, so there is an embedded transversal γ that intersects every leaf. Associated to γ there is a (complex) line bundle L over $\mathcal{F}$ whose holomorphic sections locally parameterize holomorphic functions on $\mathcal{F}$ with poles of order at most 1 at γ. We shall show that L is ample, so that $L^{\otimes k}$ has many global holomorphic sections when k is big, and the ratios of these sections give the desired map to $\mathbb{CP}^n$.

The construction of these global sections is essentially due to Poincaré. □

6. Universal circles

7. Finite depth foliations and the Thurston norm

7.1. **Surfaces and homology.** We recall some standard facts about the relation between homology classes and embedded surfaces in 3-manifolds. We work throughout with smooth manifolds and smooth maps between them.

Lemma 7.1 (Embedded surface)**.** *Let M be a compact, oriented 3-manifold. Every class in $H_2(M, \partial M; \mathbb{Z})$ is represented by the image of the fundamental class $[\Sigma]$ of an oriented compact surface Σ under a proper embedding $\Sigma \subset M$.*

Proof. Lefschetz duality says $H_2(M, \partial M; \mathbb{Z}) = H^1(M; \mathbb{Z})$. Since S^1 is a $K(\mathbb{Z}, 1)$ this latter group is in bijection with the set of homotopy classes $[M, S^1]$.

Every homotopy class of map from M to S^1 contains a smooth representative $f : M \to S^1$ for which $0 \in S^1 = \mathbb{R}/\mathbb{Z}$ is a regular value. This corresponds to a cohomology class α_f whose value on the homology class of a loop $\gamma : S^1 \to M$ is given by the winding number of $f\gamma$.

Let $\Sigma = f^{-1}(0)$. Since f is smooth and 0 is a regular value, Σ is a smooth, properly embedded surface. It is cooriented by pulling back the orientation on S^1 at 0. Since M is oriented, so is Σ, and there is a class $[\Sigma] \in H_2(M, \partial M; \mathbb{Z})$. For $\gamma : S^1 \to M$ the algebraic intersection number $\gamma \cap \Sigma$ is well-defined after we perturb γ to be in general position, and by construction this agrees with $\alpha_f([\gamma])$. Thus $[\Sigma]$ is the desired class. □

Lemma 7.2 (Regular preimage)**.** *Every compact properly embedded 2-sided surface Σ arises as the preimage of the regular value 0 for some smooth map $f : M \to S^1$ pulling back the (oriented) generator of $H^1(S^1; \mathbb{Z})$ to the class Lefschetz dual to $[\Sigma]$.*

Proof. We identify a tubular neighborhood U of Σ with $\Sigma \times (-1, 1)$. This maps by projection $\pi : \Sigma \times (-1, 1) \to (-1, 1)$ and we can pull back a form $\phi(t)dt$ by π^*, where $\phi(t)$ is a bump function with $\int_{-1}^{1} \phi(t)dt = 1$, to produce a *Thom form* θ on M. The form θ is closed, and defines (by integration) a map $f : M \to \mathbb{R}/\mathbb{Z}$ by $f(p) = \int_\gamma \theta$ where γ is any smooth path from Σ to $p \in M$. □

Corollary 7.3. *Let $A \in H_2(M, \partial M; \mathbb{Z})$ be divisible by p; i.e. $A = pB$ for some $B \in H_2(M, \partial M; \mathbb{Z})$. Then any surface Σ representing A is the disjoint union of p subsurfaces $\Sigma_1, \cdots, \Sigma_p$ each representing B.*

Proof. By Lemma 7.2 there is some $f : M \to S^1$ for which Σ is the preimage of the regular value 0. Since f pulls back the generator of $H^1(S; \mathbb{Z})$ to the class Poincaré dual to A, the image of $\pi_1(M)$ in $\pi_1(S^1)$ is contained in the subgroup of index p. Thus there is a lift $\hat{f} : M \to S^1$ so that the composition of $\hat{f}$ with the degree p cover $S^1 \to S^1$ is f. Under this cover, 0 pulls back to p distinct points in S^1, and the preimages of these points under $\hat{f}$ are disjoint surfaces $\Sigma_1, \cdots, \Sigma_p$ whose union is Σ. □

Note that we do not insist that the Σ_i are connected; and in fact if B itself is divisble, they won't be.

7.2. **Thurston norm.** Let Σ be a compact oriented surface. We denote the Euler characteristic of Σ by $\chi(\Sigma)$. If Σ is connected, define $\|\Sigma\| := \max(0, -\chi(\Sigma))$, and if the components of Σ are $\Sigma_1, \cdots, \Sigma_n$ then define

$$\|\Sigma\| = \sum \|\Sigma_i\|$$

Thus $\|\Sigma\| = -\chi(\Sigma) + 2s + d$ where s is the number of sphere components, and d is the number of disk components.

Definition 7.4 (Thurston norm)**.** Let M be a compact, oriented 3-manifold. The *Thurston norm* of a class $A \in H_2(M, \partial M; \mathbb{Z})$ is the minimum of $\|\Sigma\|$ over all surfaces representing the class A. We denote this value $\|A\|$.

The name "norm" is misleading in general, since it might take the value 0 on some nonzero class. But otherwise it satisfies the properties of a norm.

8. RFRS and the Virtual Fibration Conjecture

9. Acknowledgments

Danny Calegari was supported by NSF grant DMS 1005246.

References

[1] N. Kopell, *Commuting diffeomorphisms*, in 1970 Global Analysis (Proc. Sympos. Pure Math., Vol. XIV, Berkeley, Calif., 1968), pp. 165–184

[2] A. Navas, *Groups of Circle Diffeomorphisms*, Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 2011

[3] J. Plante, *Foliations with measure preserving holonomy*, Ann. of Math. (2) **102** (1975), no. 2, 327–361

[4] W. Thurston, *A generalization of the Reeb stability theorem*, Topology **13** (1974), 347–352

University of Chicago, Chicago, Ill 60637 USA
E-mail address: dannyc@math.uchicago.edu