CHAPTER 4: TAUT FOLIATIONS

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	1
2.	Reeb components and Novikov's Theorem	11
3.	Taut foliations	14
4.	Finite depth foliations and the Thurston norm	23
5.	Holomorphic geometry	36
6.	Universal circles	39
7.	Essential laminations	44
8.	RFRS and the Virtual Fibration Conjecture	46
9.	Acknowledgments	46
References		46

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 20, 2024.

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 $\rho : \pi_1(M) \to \text{Homeo}(F)$. 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 $\rho : \pi_1(M) \to \text{Homeo}(\mathbb{R}^n)$ 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 $\pi_1(S^1) \to \text{Homeo}(M)$ which takes the generator to φ.

1.4. Smoothness and dynamics.

The degree of smoothness of a foliation puts nontrivial constraints on the dynamics of its holonomy. This is most acute in (co)-dimension one. In this section we give three examples which are important in the theory of foliations of 3-manifolds: Kopell's Lemma, Sacksteder's Theorem (as improved by Ghys), and the Thurston Stability Theorem.

1.4.1. *Kopell's Lemma.*

Kopell's Lemma [6] 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 $g = \text{id}$.

I learned the following proof from Navas [7], 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

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

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. *Sacksteder's Theorem.* A *pseudogroup* acting on a space X is a collection of homeomorphisms between open subsets of X that may be composed and inverted on suitable intersections of their domains and ranges in the obvious sense. If $\mathcal{F}$ is a foliation and τ is a transversal, holonomy between open subsets of τ is a pseudogroup. Whether one wants to describe holonomy in terms of pseudogroups or in terms of groupoids of germs is largely a matter of taste.

Sacksteder's theorem (Thm. 1 from [11]) concerns the holonomy pseudogroup associated to a codimension 1 foliation of a compact manifold which is at least C^2. It applies when there is an *exceptional minimal set*; i.e. a closed nonempty union of leaves Λ in which every leaf is dense, and for which the intersection with a transversal is a Cantor set. It is usually stated in the language of pseudogroups.

Theorem 1.11 (Sacksteder)**.** *Let G be a finitely generated pseudogroup of C^2 diffeomorphisms of $\mathbb{R}$ leaving invariant a Cantor set K, every orbit of which is dense. Then there is some $g \in G$ fixing $x \in K$ with $g'(x) \neq 1$.*

Proof. For simplicity we prove the theorem assuming G is orientation-preserving. This hypothesis may be easily removed; we leave details to the reader. Let S be a finite set of generators for G. By restricting domains if necessary we may assume there are constants M, θ so that $M^{-1} \leq |g'(x)| \leq M$ and $|g''(x)| \leq \theta/M$ throughout the domain of any $g \in S$. Suppose we have some sequence $g_i \in S$ (not necessarily distinct) and define $h_n := g_n g_{n-1} \cdots g_1$. Suppose further that $[u, v]$ is in the domain of h_n. Then by the chain rule and the triangle inequality (as in the proof of Theorem 1.9) we have

$$|h_n'(u)| \leq |h_n'(v)| \exp\Big(\theta \sum_{j<n} |h_j(u) - h_p(v)|\Big)$$

Let's suppose K is contained in $[0, 1]$ (with endpoints at 0 and 1 for simplicity), and let's let $q \in K \cap (0, 1)$ be the supremum of some bounded open interval J_0 contained in $[0, 1] - K$. Choose an infinite sequence g_i as above so that $h_i(q) \neq h_j(q)$ for $i \neq j$. Let $q_j := h_j(q)$ be the uppermost point of a complementary interval J_j, and note that $\sum |J_j| \leq 1$. Let's further choose our sequence g_j so that some subsequence of q_n converges to q. This is possible because every orbit is dense and because G is finitely generated (so that we can pass to a diagonal subsequence). Again because of finite generation, there are only finitely many intervals J_j for which g_j is not defined on all of J_j; let's throw out this finite collection and relabel indices starting after the last one; note that we may still assume some subsequence of q_n converges to q. Thus J_0 is in the domain of h_n for all n, and therefore for each n there is some point in J_0 for which $h_n' \leq |J_n|/|J_0|$. Thus by the estimate above, $|h_n'(q)| \leq |J_n|/|J_0|e^\theta$ for all n.

Pick $p \in K$ with $|p - q| = \kappa$, and define $p_n := h_n(k)$ and $\kappa_n := |p_n - q_n|$. Evidently

$$\kappa_n \leq \kappa h_n'(q) \exp\Big(\theta \sum_{i<n} \kappa_i\Big) \leq \kappa |J_n|/|J_0| \exp\Big(\theta(1 + \sum_{i<n} \kappa_i)\Big) \tag{1.1}$$

providing $[q, p]$ is in the domain of h_n.

We shall prove by induction that $[q, p]$ is in the domain of h_n and that $\kappa_n \leq C|J_n|$ for any fixed C providing κ is chosen sufficiently small. By finite generation there is a positive constant ϵ so that if g_i is any generator, and $x \in K$ is in the domain of g_i, then any $y \in K$

with $|x-y|<\epsilon$ is also in the domain of g_i. Thus to prove the induction hypothesis it will suffice to prove the estimate on κ_n (taking C small compared to ϵ).

The induction hypothesis implies $\sum_{j<n}\kappa_j \le C\sum_{j<n}|J_j| \le C$. Thus, assuming $\kappa \le C|J_0|$ and using Equation 1.1 we obtain $\kappa_j \le C|J_j|\exp(\theta(1+C))$ for all $j<n$ and therefore $\sum_{j=1}^{n-1}\kappa_j \le C\exp(\theta(1+C))$. Substituting this in Equation 1.1 again gives the estimate

$$\kappa_n \le \kappa|J_n|/|J_0|\exp\Big(\theta(1+Ce^{\theta(1+C)})\Big)$$

so providing $\kappa \le C|J_0|\exp(-\theta(1+Ce^{\theta(1+C)}))$ the induction step is proved.

Now observe that for n big, $h_n' \sim |J_n| \ll 1$ for big n on all of $[q,p]$. If we choose some n for which q_n is sufficiently close to q then h_n will take some interval I about q properly inside itself; thus there will be a fixed point (necessarily in K) and at this fixed point $h_n' < 1$. □

The following observation is due to Ghys (see [3], Lem. 1.2.9 whose argument we reproduce essentially verbatim).

Lemma 1.12 (Ghys)**.** *Let G be a pseudogroup of homeomorphisms of $\mathbb{R}$ with dense orbits. Then either G has trivial holonomy or there is a finitely generated subpseudogroup G_0 with an exceptional minimal set.*

Proof. Again, we assume G is orientation-preserving for simplicity. If there is non-trivial holonomy then there is some interval $I := [a,b]$ and some element $f \in G$ with $f(a)=a$ and $f(x) < x$ for all $a < x \le b$. Since orbits are dense there is g defined on some smaller interval $[a,a']$ with $a < g(a) < b$. Choosing a' small enough we may further assume $a' < g(a) < g(a') < b$. Replacing f by a power if necessary we may assume $a < f(b) < a' < g(a) < g(f(b)) < b$. Let $f_1 = f$ and $f_2 = gf$. Then both f_1 and f_2 take I to disjoint intervals $I_j := f_j(I)$ inside I, and therefore the semigroup S they generate leaves invariant a Cantor set

$$K = \bigcap_n \bigcup_{w\in S,\ |w|=n} w(I)$$

□

Combining Lemma 1.12 and Theorem 1.11 and restating in the language of foliations gives:

Theorem 1.13 (Linear holonomy)**.** *Let $\mathcal{F}$ be a C^2 codimension 1 foliation of a compact manifold M. Then*

(1) every exceptional set has nontrivial linear holonomy; and
(2) if $\mathcal{F}$ is minimal, either $\mathcal{F}$ has no holonomy at all or $\mathcal{F}$ has nontrivial linear holonomy.

Codimension one foliations with no holonomy at all are very special, particularly in 3-manifolds. We shall return to this subject in § 3.2. For now we restrict ourselves to the following observation, due to Hölder:

Proposition 1.14 (Hölder)**.** *Suppose G is a nontrivial group of homeomorphisms of $\mathbb{R}$ with no holonomy. Then G is (free) abelian, isomorphic to a subgroup of $\mathbb{R}$, and the action is semi-conjugate to a group of translations.*

Proof. Since there is no holonomy, it follows that no nontrivial element of G fixes any point in $\mathbb{R}$. It follows that for each nontrivial g either $g(x) > x$ for all x or $g(x) < x$ for all x. This implies in particular for any nontrivial g and for any $x > 0$ that there is an integer n (possibly negative) with $g^n(0) > x$, or else g would fix $\sup_n g^n(0)$ (this is one definition of the so-called *Archimedean property* for G). Now, let $g \in G$ be arbitrary with $g(0) > 0$ and then for any $h \in G$ observe that there is a Dedekind cut (L_h, U_h) consisting of all rational numbers p/q for which $g^p(0) \leq h^q(0)$ and for which $g^p(0) > h^q(0)$ respectively. By the Archimedean property one may verify the following facts:

(1) both L_h and U_h are nonempty;
(2) every element of L_h is less than every element of U_h;
(3) if $\rho(h) \in \mathbb{R}$ is the real number associated to the Dedekind cut (L_h, U_h) then ρ is an injective homomorphism.

Now let $G0$ denote the orbit of 0. If $G0$ is discrete it has no accumulation points, and is order isomorphic to $\mathbb{Z} \subset \mathbb{R}$. In this case G is isomorphic to $\mathbb{Z}$ and the action is semi-conjugate to a group of (integer) translations. Otherwise G is isomorphic to a dense subgroup $\rho(G)$ of $\mathbb{R}$. This group acts transitively and in an order preserving manner on the orbit $G0$. Take the closure $\overline{G0}$ and collapse complementary gaps to points to obtain a semi-conjugate action which by construction is topologically isomorphic to the action of $\rho(G)$ on $\mathbb{R}$ by translations. □

1.4.3. *Thurston Stability Theorem.* The Thurston's Stability Theorem [15] applies in any (co)-dimension. It is concerned with the kernel of the linear part of holonomy. Unlike the theorems of Kopell and Sacksteder it applies when the action is only C^1.

Theorem 1.15 (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 $G \to \mathrm{GL}(n, \mathbb{R})$ 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 $\le k$ (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 $\text{Hom}(G; \mathbb{R}) \neq 0$.

Example 1.16. 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 $\rho : \pi_1(\lambda) \to \text{Homeo}(S^1)$ be the induced action on the circle at infinity.

Since $\pi_1(\lambda)$ is free, we can lift ρ to a representation $\tilde{\rho} : \pi_1(\lambda) \to \text{Homeo}(\mathbb{R})$ 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 $\sigma : \pi_1(\lambda) \to \text{Homeo}(I)$, 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 1.17 (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.

□

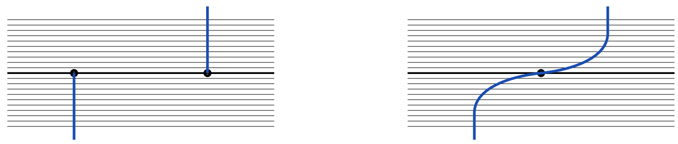

FIGURE 1. A nonclosed leaf admits a transversal

The next Theorem is due to Novikov [8], and is a cornerstone of the codimension 1 theory.

Theorem 1.18 (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 1.17 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 1.19. 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 1.20. *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 1.21 (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 1.18 and Lemma 1.20 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 1.21 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 $z \ge 0$. 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 § **??**.

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. *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. *There is an immersed circle $\gamma : S^1 \to M$ transverse to $\mathcal{F}$ and intersecting every leaf of $\mathcal{F}$.*
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. *There is a closed 2-form ω on M^3 positive on $T\mathcal{F}$.*

(5) *There is a flow X transverse to $\mathcal{F}$ which is volume preserving for some Riemannian metric on M.*
(6) *There is a Riemannian metric on M for which every leaf of $\mathcal{F}$ is a minimal surface.*
(7) *There is a Riemannian metric on M and a closed 2-form ω which calibrates $\mathcal{F}$.*
(8) *There is a map $f : M \to S^2$ whose restriction to each leaf λ is a branched covering $f : \lambda \to S^2$.*
(9) *For any Riemannian metric on M, there is a map $f : M \to \mathbb{CP}^n$ for some $n \geq 1$ which is leafwise holomorphic (in the induced conformal structure).*
(10) *For any Riemannian metric on M, there is a symplectic W^4, ω and almost-complex structure J and a map $f : M \to W$ which is leafwise (pseudo)-holomorphic.*

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 1.21 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; the equivalence of the last condition with the others will be deferred until Chapter 5.

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.

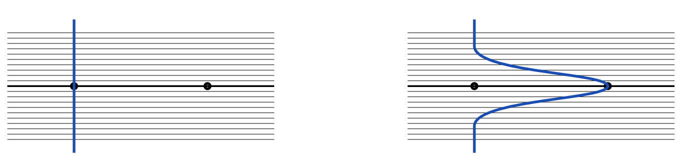

FIGURE 2. Leafwise homotop γ so it passes through q

□

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

Proof. By Proposition 1.17 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 $\omega(\xi) \le \text{area}(\xi)$ 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

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

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

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

□

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 proves the equivalence of conditions (1)–(7) in 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 [10]:

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. subsurfaces such that

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

The boundaries $\partial\lambda_i$ might *a priori* be very wiggly; if so, since M is compact (and therefore leaves of $\mathcal{F}$ have bounded geometry), we may adjust the $\partial\lambda_i$ in λ to smooth hypersurfaces with bounded geometry and no greater volume, at the cost of adjusting $\text{volume}(\lambda_i)$ by a term comparable to volume($\partial\lambda_i$); thus we may assume the $\partial\lambda_i$ have bounded geometry. 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 homotopy can be perturbed to intersect $\partial\lambda_i$ a bounded number of times (independent of i), and the Følner condition then guarantees $|\mu_i(\tau') - \mu_i(\tau)| \to 0$. No transversal τ can intersect any λ_i in more than $C(\tau)$volume(λ_i) points, and since M is compact, a total transversal will intersect λ_i for big i in $\epsilon \cdot$ volume(λ_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 λ_i/volume(λ_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 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.

We now give two conditions that imply the existence of a nontrivial invariant measure. The first is in terms of the conformal geometry of leaves, and the second is in terms of holonomy.

3.2.1. *Parabolic leaves.* Let $\mathcal{F}$ be a foliation of a 3-manifold M. A Riemannian metric on M determines a conformal structure on the leaves of $\mathcal{F}$, and a leaf λ is *parabolic* if its universal cover is conformally isomorphic to $\mathbb{C}$. This property of a leaf does not depend on a choice of a metric, since different metrics induce quasiconformally equivalent conformal structures, and any Riemann surface quasiconformally isomorphic to $\mathbb{C}$ is actually isomorphic to $\mathbb{C}$ by the measurable Riemann mapping theorem.

The following is due to Candel [2]:

Proposition 3.14 (Parabolic leaf)**.** *Let λ be a parabolic leaf. Then there is a nontrivial invariant measure μ on $\mathcal{F}$ with support contained in the closure of λ.*

Proof. As in the proof of Proposition 3.10 it suffices to find a sequence of subsets λ_i of λ with $\text{length}(\partial\lambda_i)/\text{area}(\lambda_i) \to 0$. It is evidently enough to find such a sequence in the universal cover of λ, so without loss of generality we fix an isomorphism $f : \mathbb{C} \to \lambda$, let B_t denote the ball of radius t, and set $\lambda_t = f(B_t)$. Then

$$\text{area}(\lambda_t) = \int_{B_t} |df|^2 = \int_{r=0}^{t} \left(\int_{\partial B_r} |df|^2 \right) dr$$

On the other hand,

$$\text{length}(\partial\lambda_t) = \int_{\partial B_t} |df|$$

so by Cauchy–Schwarz,

$$\text{length}(\partial\lambda_t)^2 = \left(\int_{\partial B_t} |df| \right)^2 \le 2\pi t \int_{\partial B_t} |df|^2 = 2\pi t \frac{d}{dt}\text{area}(\lambda_t)$$

Now, for any fixed t

$$\int_t^\infty \frac{\text{area}'(\lambda_r)}{\text{area}(\lambda_r)^2} dr = \frac{1}{\text{area}(\lambda_t)} < \infty$$

On the other hand, if $\liminf \text{length}(\partial\lambda_t)/\text{area}(\lambda_t) > 0$ then there is a positive constant C so that

$$\int_t^\infty \frac{\text{area}'(\lambda_r)}{\text{area}(\lambda_r)^2} dr \ge C \int_t^\infty \frac{\text{area}'(\lambda_r)}{\text{length}(\partial\lambda_r)^2} dr \ge C \int_t^\infty \frac{dr}{2\pi r} = \infty$$

This proves the proposition. □

3.2.2. *Branching and holonomy.* Let $\mathcal{F}$ be a taut foliation of a 3-manifold M. Let $\tilde{\mathcal{F}}$ be the foliation of the universal cover $\tilde{M}$. By Theorem 2.4 the leaf space L of $\tilde{\mathcal{F}}$ is a simply-connected 1-manifold. However, it is entirely possible for L to fail to be Hausdorff: there might be distinct leaves μ, μ' of $\tilde{\mathcal{F}}$ and a sequence of leaves λ_i of $\tilde{\mathcal{F}}$ with points $p_i, q_i \in \lambda_i$ so that $p_i \to p \in \mu$ and $q_i \to q \in \mu'$. We refer to this phenomenon as *branching of the leaf space*; if L is Hausdorff (and therefore homeomorphic to $\mathbb{R}$) we say $\mathcal{F}$ is $\mathbb{R}$-*covered.*

First we prove a proposition.

Proposition 3.15 (Branch implies holonomy)**.** *Let $\mathcal{F}$ be a foliation of a 3-manifold M and suppose that $\mathcal{F}$ is not $\mathbb{R}$-covered. Then $\mathcal{F}$ has nontrivial holonomy.*

Proof. Let's suppose $\mathcal{F}$ has no holonomy. If $\mathcal{F}$ contains a closed surface S, then an open neighborhood of S is foliated as a product by copies of S. By Theorem 1.18 the set of closed leaves is closed, so if it is open it is all of M and therefore M is an S-bundle over S^1 which is certainly $\mathbb{R}$-covered (notice by the way this implies that a foliation with no holonomy is taut).

If $\mathcal{F}$ contains an exceptional minimal set Λ we let N be a component of $M - \Lambda$, and let $\overline{N}$ be the metric completion of N in the path metric. We may split N into two parts — a collection of disjoint foliated I-bundles over non-compact surfaces, each with a single boundary component (which, by the lack of holonomy, are necessarily foliated as a product) and a compact codimension 0 submanifold G; the boundary of G decomposes into the *horizontal boundary* $\partial_h G$ contained in leaves of Λ, and the *vertical boundary* $\partial_v G$ consisting of trivially foliated annuli transverse to $\mathcal{F}$ that cut off the noncompact I-bundles. Let F be a component of $\partial_h G$; note that F is a compact surface with boundary. Since there is no holonomy, an open neighborhood of F in G is foliated as a product by copies of F. Once again, Theorem 1.18 implies that the set of closed leaves of $\mathcal{F}|G$ is closed, and therefore as before $G = F \times I$ foliated by copies of F. Gluing back on the noncompact I-bundles, we see that every complementary region to $\mathcal{F}$ is an I-bundle foliated as a product; we may collapse these I-bundles (this is the inverse of the Denjoying operation) to obtain a new foliation $\mathcal{F}'$ which is $\mathbb{R}$-covered if and only if $\mathcal{F}$ is, and has nontrivial holonomy if and only if $\mathcal{F}$ does.

Thus we are reduced to the case that $\mathcal{F}$ is minimal. Suppose the leaf space L of $\tilde{\mathcal{F}}$ branches. Since $\mathcal{F}$ is minimal there are ϵ and t so that there is λ_i, μ, μ' and $p_i, q_i \in \lambda_i$ with $p_i \to p \in \mu$ and $q_i \to q \in \mu'$. Without loss of generality we may suppose $p_i \to p$ and $q_i \to q$ from below. Let ϵ be such that $\mathcal{F}$ may be covered by product charts in such a way that every metric ball of radius ϵ is contained in a product chart. By minimality there is a number $t > 0$ such that for any metric ball U of radius $\epsilon/3$ every point in $\mathcal{F}$ may be joined by a leafwise path of length less than t. On the other hand, by compactness, there is a number $\delta > 0$ so that if τ is any transveral of length $\leq \delta$ then holonomy transport along any path of length $\leq t$ takes τ to a path of length $\leq \epsilon/3$. For each i let γ_i be a path in λ_i from p_i to q_i. If the pocket of leaves between λ_i and μ were contained in product charts all along γ_i, then we would obtain a leafwise path from μ to μ' which is absurd; thus this pocket must get large at some point, and there is an ϵ-ball U_i contained in it. Let V_i be the $\epsilon/3$-ball with the same center as U_i. Let i be such that p_i and p are contained in a transversal τ of length $< \delta$ and let δ be a path in λ_i of length $\leq t$ from p_i to some point r for which there is $\alpha \in \pi_1$ with $\alpha(r) \in V_i$. Then holonomy transport of $\alpha(\tau)$ along $\alpha(\delta)$

takes it entirely inside U_i. If we let I denote the interval in the leafspace L of $\tilde{\mathcal{F}}$ that τ projects to, then α takes I properly inside itself, so that it necessarily has a fixed point corresponding to a leaf of $\tilde{\mathcal{F}}$ that covers a leaf of $\mathcal{F}$ with nontrivial holonomy. □

Remark 3.16. In the proof of Proposition 3.15 we considered the case that $\mathcal{F}$ has an exceptional mininal set Λ, and considered a decomposition of the path completion $\overline{N}$ of a complementary region N. In this decomposition the submanifold G is called the *guts* and the noncompact I-bundles are the *interstitial regions*. We shall meet such decompositions again in § 7.

From this we may deduce:

Proposition 3.17. *Let $\mathcal{F}$ be a foliation without holonomy of a 3-manifold M. Then every minimal set of $\mathcal{F}$ admits an invariant transverse measure of full support. Furthermore, M is a surface bundle over a circle, and after possibly collapsing some complementary I-bundles foliated as a product, $\mathcal{F}$ may be perturbed to a foliation by closed surfaces.*

Proof. We have already seen that if $\mathcal{F}$ has a closed surface then M is a surface bundle over S^1 and $\mathcal{F}$ is a foliation by closed surfaces, in which case we are done. Likewise if $\mathcal{F}$ has an exceptional minimal set we have already seen that we may collapse complementary regions (which are I-bundles) to obtain a new foliation that is minimal, $\mathbb{R}$-covered (by Proposition 3.15) and has no holonomy.

Thus $L = \mathbb{R}$ and since $\mathcal{F}$ has no holonomy, by Proposition 1.14 the action of $\pi_1(M)$ on L is conjugate to an action by translations. It follows that $\mathcal{F}$ admits a transverse measure of full support and without atoms. Disintegrating this measure in a product chart gives a closed nowhere zero 1-form α with $\ker(\alpha) = T\mathcal{F}$. This 1-form may be perturbed to have rational periods (staying nowhere zero) exhibiting M as a surface bundle. □

4. Finite depth foliations and the Thurston norm

4.1. 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 4.1 (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 $[S]$ of an oriented compact surface S under a proper embedding $S \subset M$.*

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 $S = f^{-1}(0)$. Since f is smooth and 0 is a regular value, S 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 S, and there is a class $[S] \in H_2(M, \partial M; \mathbb{Z})$. 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 $[S]$ is the desired class. □

Lemma 4.2 (Regular preimage)**.** *Every compact properly embedded 2-sided surface S 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 $[S]$.*

Proof. We identify a tubular neighborhood U of S with $S\times(-1,1)$. This maps by projection $\pi : S \times (-1,1) \to (-1,1)$ 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 S to $p \in M$. ∎

Corollary 4.3. *Let $\alpha \in H_2(M, \partial M;\mathbb{Z})$ be divisible by p; i.e. $\alpha = p\alpha'$ for some $\alpha' \in H_2(M, \partial M;\mathbb{Z})$. Then any surface S representing A is the disjoint union of p subsurfaces $S_1, \cdots, S_p$ each representing α'.*

Proof. By Lemma 4.2 there is some $f : M \to S^1$ for which S 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 α, 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 $\S_1, \cdots, S_p$ whose union is S. ∎

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

4.2. **Thurston norm.** Let S be a compact oriented surface. We denote the Euler characteristic of S by $\chi(S)$. If S is connected, define $\|S\| := \max(0, -\chi(S))$, and if the components of S are $S_1, \cdots, S_n$ then define

$$\|S\| = \sum \|S_i\|$$

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

Definition 4.4 (Thurston norm)**.** Let M be a compact oriented 3-manifold. The *Thurston norm* of a class $\alpha \in H_2(M, \partial M;\mathbb{Z})$ is the minimum of $\|S\|$ over all surfaces representing the class α. We denote this value $\|\alpha\|$.

The name "norm" is misleading in general, since it might take the value 0 on some nonzero class. Recall that a function $\|\cdot\|$ from a real vector space to $\mathbb{R}$ is called a *semi-norm* if it is convex, non-negative, even, and linear on rays.

Proposition 4.5. *Suppose M is irreducible and ∂M is incompressible. The function $\|\cdot\|$ extends uniquely to a seminorm on $H_2(M, \partial M;\mathbb{R})$.*

Proof. The function $\|\cdot\|$ is evidently even and non-negative. It is linear on rational rays by Corollary 4.3. Therefore it suffices to show that it is convex. Let S, S' be two embedded oriented surfaces with no sphere or disk component representing integral homology classes α and α'. Put them in general position, and by an isotopy eliminate innermost curves or arcs of intersection that are inessential in either surface. Let S'' be obtained from $S \cup S'$ by resolving the intersections to produce a new embedded oriented surface. Because no intersection was inessential in either surface, no component of S'' is a disk or a sphere. Thus $\|S''\| = \|S\| + \|S'\|$ and we are done. □

Theorem 4.6 (Thurston norm)**.** *If M has no homologically essential torus or annulus, the seminorm $\|\cdot\|$ is a genuine norm. The unit ball for $\|\cdot\|$ is a finite sided rational polyhedron.*

Proof. Choose a basis $b_1, \cdots, b_n$ for $H_2(M, \partial M; \mathbb{Z})$ and identify $H_2(M, \partial M; \mathbb{R})$ with $\mathbb{R}^n$ for some n. Choose any representative surfaces S_i for the basis elements and let $T = \max \|S_i\|$. By convexity and evenness, $\|\cdot\| \le T\|\cdot\|_1$ where $\|\cdot\|_1$ denotes the L^1 norm in the given basis. Thus the unit ball for $\|\cdot\|$ contains an open neighborhood U of 0.

Let $\mathcal{V}$ denote the set of linear functions $v : \mathbb{R}^n \to \mathbb{R}$ taking integer values on $\mathbb{Z}^n$, and let $\mathcal{P}$ denote the set of hyperplanes where $v = 1$ for some $v \in \mathcal{V}$. Observe that there are only finitely many elements of $\mathcal{P}$ that do not intersect U.

If b'_j is any other basis, there is a unique linear function $v \in \mathcal{V}$ agreeing with $\|\cdot\|$ on the b'_j. Let $\pi(v) \in \mathcal{P}$ be the hyperplane where $v = 1$. Note that if w is not a positive linear combination of b'_j then $\|w\| \ge v(w)$ by convexity, and therefore π does not intersect U except possibly in the cone spanned by the b'_j.

If r is any projective ray with $\|r\| > 0$, we may find a sequence of bases $b_{j,n}$ that converge projectively to r (so that r is in the positive cone for each basis), and let $v_n \in \mathcal{V}$ agree with $\|\cdot\|$ on the $b_{j,n}$ with hyperplane $\pi(v_n)$. When n is large, the intersection of $\pi(v_n)$ with the cone on the $b_{j,n}$ is close to the boundary of the unit ball; in particular, it lies outside U. But then $\pi(v_n)$ lies outside U for all large n, so that v_n is eventually constant and agrees with $\|\cdot\|$ in a projective neighborhood of r. It follows that the boundary of the unit ball is cut out by some subset of the finitely many $\pi(v)$ that do not intersect U, so it is a finite sided rational polyhedron. Note that this argument works even if $\|\cdot\|$ is degenerate somewhere.

It follows that $\|\cdot\|$ is piecewise rational linear, so if it takes the value 0 on some nonzero vector, it takes the value 0 on a rational vector. Such a vector projectively represents a nonzero integral class α with $\|\alpha\| = 0$. Note that the only properties of $\|\cdot\|$ used up to this point are that it is a seminorm taking integer values on integer vectors.

A norm-minimizing surface projectively representing a rational class with norm 0 is necessarily an essential torus or annulus. □

If M has no homologically essential torus or annulus, there is a dual norm on the vector space $H^2(M, \partial M; \mathbb{R})$ given by the natural pairing with H_2. The unit ball in the dual norm is likewise a finite sided rational polyhedron, whose vertices lie in $H^2(M, \partial M; \mathbb{Z})$.

Remark 4.7. If M is compact and irreducible and $F \subset \partial M$ is a closed incompressible subsurface, we may define the Thurston norm on $H_2(M, F; \mathbb{Z})$ in the obvious way, and extend it to a seminorm on $H_2(M, F; \mathbb{R})$ which will be a norm unless some class is represented by a surface with $\chi = 0$. The unit ball will be a finite sided rational polyhedron; the proof goes through as above with essentially no modification.

4.3. **Euler class inequality.**

Proposition 4.8. *Let $\mathcal{F}$ be a taut foliation, let S be any closed oriented surface, and let $f : S \to M$ be any map. Then $|f^*e(\mathcal{F})[S]| \le \|S\|$.*

Proof. Changing the coorientation of $\mathcal{F}$ changes $e(\mathcal{F})$ to $-e(\mathcal{F})$ so it suffices to show $f^*e(\mathcal{F})[S] \le \|S\|$.

By abuse of notation we identify S with its image, which we take to be an immersed surface in M. If S is compressible we may compress it without changing its homology class or increasing $\|S\|$ until it is incompressible. Choose a metric on M for which leaves of $\mathcal{F}$ are minimal surfaces, and homotop S to a minimal representative. Then either S is homotopic to a cover of a leaf of $\mathcal{F}$ in which case $e(\mathcal{F})[S] = \pm\chi(S)$, or the foliation $\mathcal{F} \cap S$ has only finitely many singularities at the (isolated) tangencies of S with leaves of $\mathcal{F}$. Each tangency s is a saddle or generalized saddle; at each of them the index of the singular foliation index(s) is a negative integer. We may compute $\chi(S)$ by summing these indices. On the other hand, we may compute $e(\mathcal{F})[S]$ by summing these indices times a character which is ± 1 depending on whether the co-orientations of S and $\mathcal{F}$ agree or disagree. Thus

$$e(\mathcal{F})[S] = \sum_s \pm\,\text{index}(s) \le -\sum_s \text{index}(s) = -\chi(S) = \|S\|$$

□

Corollary 4.9. *Suppose S is a compact leaf of a taut foliation $\mathcal{F}$. Then S is norm minimizing and $e(\mathcal{F})$ is in the boundary of the unit ball of the dual Thurston norm.*

Proof. This follows from the equality $e(\mathcal{F})[S] = \chi(S)$ together with Proposition 4.8. □

The remainder of this section is devoted to proving a kind of converse to this, due to Gabai:

Theorem 4.10 (Gabai, minimizer is leaf)**.** *Let M be compact, oriented, and irreducible with incompressible boundary. For every primitive nonzero class $\alpha \in H_2(M, \partial M; \mathbb{Z})$ and every embedded surface S representing α of minimal norm there is a taut co-oriented foliation $\mathcal{F}$ of M with S as a leaf.*

It follows that every vertex of the dual Thurston norm ball is the Euler class of some taut foliation.

The contents of the next few sections are all essentially due to Gabai [4] and we follow his paper with only minor modifications.

4.4. **Sutured manifolds.** Theorem 4.10 is proved by induction with the help of an auxiliary combinatorial structure called a *taut sutured hierarchy*, analogous in many ways to a hierarchy for a Haken manifold. The terms in the hierarchy are sutured manifolds.

In the (usual) theory of hierarchies one obtains manifolds with corners (sometimes expressed using the terminology 'boundary pattern') where facets of successive cutting surfaces meet; keeping track of this combinatorial data is important at many points in the theory and its applications. For example, it is the key to the orbifold trick that reduces the inductive gluing step in Thurston's hyperbolization theorem to the last step. In Gabai's definition of a sutured manifold [4] the boundary of the manifold, and the boundary of proper surfaces in it, are smooth; this leads to formulae for (relative) Thurston norm, Euler characteristic etc. that fails to satisfy the natural properties we would like (additivity, convexity), and the correct treatment of disks leads to a proliferation of special cases. Scharlemann [12] worked out a modified version of the theory with better properties and that works in some generality. We use the language of manifolds with corners, which is intermediate between Gabai's theory and Scharlemann's.

Definition 4.11. A *sutured manifold* is a triple (M, γ, β) where M is a compact oriented 3-manifold, the *sutures* $\gamma = A(\gamma) \cup T(\gamma)$ consist of a union of disjoint closed annuli and tori in ∂M, and β is a family of essential arcs dividing some annular sutures into rectangles. Thus the trivalent graph $\Gamma := \partial A(\gamma) \cup \beta$ gives M the structure of a manifold with corners.

If $R(\gamma)$ denotes the closure of $\partial M - \gamma$, the components of $R(\gamma)$ are oriented R_+ and R_- (depending on whether the orientation agrees or disagrees with that of ∂M) in such a way that for each annulus A in $A(\gamma)$ one boundary component is on ∂R_+ and one is on ∂R_-.

Note that for each annulus A in $A(\gamma)$ the two components of ∂A inherit an orientation from $\partial R_\pm$ that agree after an isotopy across A. This may be expressed by saying that the core of each annulus gets a canonical orientation, agreeing the with the induced orientations on ∂A up to isotopy. Each component of β may also be oriented so that it runs from R_- to R_+; with this convention, the intersecton number of a component of ∂A with a component of β is positive.

Example 4.12 (Product sutured manifold)*.* Let F be a compact and oriented (not necessarily connected) surface with corners $V \subset \partial F$. Then $(F \times I, \partial F \times I, V \times I)$ is a sutured manifold (called a *product* sutured manifold) in an obvious way with $\gamma = A(\gamma)$ and $R_+ = F \times 1$ and $R_- = F \times 0$.

Let (M, γ, β) be a sutured manifold, and let $S \subset M$ be a properly embedded surface. Then S inherits the structure of a surface with corners, and we may define $\chi(S)$ by subtracting $1/4$ from the usual Euler characteristic at each corner. Let $F \subset \partial M$ be some subsurface. We may define the Thurston norm $\|\cdot\|$ on $H_2(M, F)$ with respect to this modified Euler characteristic, discarding surfaces with $\chi > 0$ (note that this takes non-negative values in $\frac{1}{4}\mathbb{Z}$). Providing there are no essential surfaces in M with $\chi > 0$ Proposition 4.5 and Theorem 4.6 go through essentially verbatim: $\|\cdot\|$ extends to a seminorm, which is an honest norm unless M contains an essential surface with $\chi = 0$, and whose unit ball is a finite sided rational polyhedron (possibly noncompact).

Definition 4.13. A sutured manifold (M, γ, β) is *taut* if M is irreducible and both $R_\pm$ are incompressible and norm minimizing in $H_2(M, \gamma)$.

Example 4.14*.* A product sutured manifold is taut if and only if $\chi(F) \leq 0$ as a surface with corners.

Example 4.15 (Taut sutured handlebody)*.* Figure 3 shows a taut sutured genus 2 handlebody H, γ (in this example β is empty). There are three annular sutures that decompose ∂H into two thrice-punctured spheres. Since the core of the annuli are essential and nonparallel in H, these thrice-punctured spheres are norm minimizing, so this sutured manifold is taut.

Definition 4.16 (Decomposing surface)**.** Let (M, γ, β) be a sutured manifold, and let S be a compact properly embedded oriented surface. Suppose one of the following holds for each component λ of $\partial S \cap \gamma$:

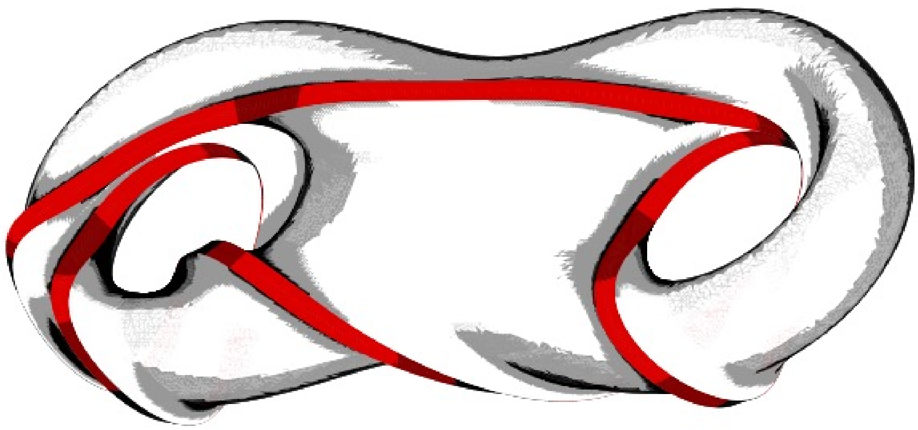

FIGURE 3. A taut sutured handlebody.

(1) λ is a proper essential arc in $A(\gamma)$ so that all intersection points of $\lambda \cap \beta$ have positive sign;
(2) λ is isotopic to, and coherently oriented with the core of a component of $A(\gamma)$ and meets β essentially (equivalently, all intersection points of $\lambda \cap \beta$ have positive sign); or
(3) λ is an essential curve in some component T of $T(\gamma)$, and every component of $\partial S \cap T$ is parallel to, and coherently oriented with λ.

Suppose further that no circle component of $\partial S \cap R$ bounds a disk in either S or R, and that no arc of $\partial S \cap R$ is inessential in R. Then S is a *decomposing surface* for (M, γ, β).

Let (M, γ, β) be a sutured manifold, and let S be a decomposing surface. Let $M' = M - N(S)$ where $N(S)$ is a neighborhood of S. Let $S'_\pm$ be the copies of S in $\partial M'$, where S'_+ (resp. S'_-) is the copy of S for which the co-orientation of S points out of (resp. into) M'.

Then M' inherits the structure of a sutured manifold as follows:

(1) the sutures γ' are unions
$$\gamma' = (\gamma \cap M') \cup N(\partial S'_+ \cap R_-) \cup N(\partial S'_- \cap R_+)$$
(2) the curves β' are the closures of the components of $\beta - \partial S$ in M' together with the arcs $N(\partial S'_+ \cap R_-) \cap (\gamma \cap M')$ and $N(\partial S'_- \cap R_+) \cap (\gamma \cap M')$.

We write $(M, \gamma, \beta) \overset{S}{\rightsquigarrow} (M', \gamma', \beta')$ and say that the decomposition is taut if (M, γ, β) and (M', γ', β') are taut; see Figure 4.

Definition 4.17 (Taut sutured hierarchy)**.** A *taut sutured hierarchy* is a finite sequence of taut decompositions
$$(M_0, \gamma_0, \beta_0) \overset{S_1}{\rightsquigarrow} (M_1, \gamma_1, \beta_1) \overset{S_2}{\rightsquigarrow} \cdots \overset{S_n}{\rightsquigarrow} (M_n, \gamma_n, \beta_n)$$
where (M_n, γ_n, β_n) is a product sutured manifold.

Our goal is to prove:

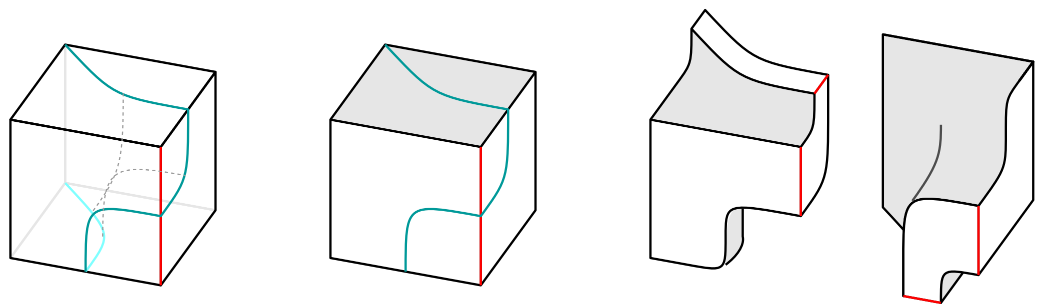

FIGURE 4. How γ and β in ∂M decompose after cutting along S into γ' and β' in $\partial M'$. The surfaces R and R' are in gray, the sutures γ and γ' are in white, the curve ∂S is in blue, and the arcs β and β' are in red.

Theorem 4.18 (Gabai; taut sutured hierarchy exists)**.** *Any taut sutured manifold admits a taut sutured hierarchy.*

Following Gabai we shall prove this in two steps. First: any taut sutured manifold which is not a product admits a taut decomposition of a special kind; and second: there is a well-ordered complexity on taut sutured manifolds that vanishes only on product sutured manifolds, and that strictly decreases along taut decompositions of the special kind.

Definition 4.19 (Special decomposing surface)**.** A decomposing surface S for (M, γ, β) is *special* if either

(1) S is connected and $\partial S \subset \gamma$ but S is not isotopic to a component of R; or
(2) ∂S intersects each nonplanar component of R in a (possibly empty) family of parallel homologically essential loops, and each planar component in a (possibly empty) family of parallel homologically essential arcs.

Lemma 4.20. *Let (M, γ, β) be a taut sutured manifold, and suppose some component of ∂M is not a sphere. Then a special decomposing surface exists.*

Proof. We find S in a component of M whose boundary is not a sphere. For any compact oriented 3-manifold the kernel of $H_1(\partial M) \to H_1(M)$ (coefficients in $\mathbb{R}$) is a Lagrangian subspace L. Let J be the subspace of $H_1(\partial M)$ spanned by the cores of $A(\gamma)$ (note that J is isotropic), and let J_R be the subspace of J spanned by ∂R_i where R_i ranges over the components of R.

Certainly $J_R \subset J \cap L$. If $J \cap L$ is bigger than J_R it follows that $H_2(M, \gamma)$ is not generated by the $[R_i]$, and therefore we may find a primitive integral class α not in the image. Add sufficiently many copies of $[R_+]$ to α if necessary so that $\partial\alpha$ is a non-negative multiple of cores of $A(\gamma)$; then any norm-minimizing surface S representing α will be a special decomposing surface.

So let's assume $J_R = J \cap L$. The dimension of J is equal to half the dimension of $H_1(\partial M)$ if and only if R is planar and $T(\gamma)$ is empty. Otherwise intersection number defines a map $L/J_R \to (J/J_R)^*$ which necessarily has a nontrivial kernel. In this case let $z \in H_1(\partial M; \mathbb{Z}) \cap L$ be nonzero and have trivial intersection number with each component of $A(\gamma)$. By adding a sufficiently large multiple of $[\partial R_+]$ if necessary, we may assume z

is represented by a 1-manifold Γ whose components either lie in $A(\gamma)$ (and are positively oriented) or are disjoint from $A(\gamma)$ and essential in R. For each component R_i of R we may replace $\Gamma \cap R_i$ by a homologous collection of parallel essential loops, and then any norm-minimizing surface S representing a non-trivial class in $H_2(M, \Gamma)$ will be special.

Finally we are in the case that every component of R is planar, $T(\gamma)$ is empty, and $J \cap L = J_R$. Thus $L/J_R \to (J/J_R)^*$ is surjective, so we may choose a nonzero class $z \in H_1(\partial M; \mathbb{Z})$ with the property that for each component R_i, intersection number with z is nonzero for at most two components of ∂R_i. We may thus represent z by Γ whose intersection with each R_i is a family of parallel essential arcs, and choose a norm-minimizing surface S representing a nontrivial class in $H_2(M, \Gamma)$ as above. ☐

Lemma 4.21. *Let (M, γ, β) be taut. Then there is a special decomposing surface S that induces a taut decomposition $(M, \gamma, \beta) \overset{S}{\rightsquigarrow} (M', \gamma', \beta')$.*

Proof. Let S be a special decomposing surface, representing some class in $H_2(M, \partial S \cup \gamma)$. Because the unit ball of the Thurston norm is a finite sided polyhedron, there is a positive n_0 so that the classes $[S] + n[R]$ and $[R]$ intersect the same (closed) face of the norm ball for all $n \geq n_0$. Replace S by a new norm minimizing surface S' representing the class of $[S] + n[R]$. The surface S' is special if S is. Furthermore, both S' and $S' + R$ are norm minimizing in M and therefore also in M'. But $S' + R = R'$ so R' is norm-minimizing in M'. Evidently M' is irreducible if M is, so M' is taut. ☐

Certain surfaces always induce taut decompositions:

Lemma 4.22. *Let (M, γ, β) be taut, and let S be a decomposing surface which is either of the following:*

1. *an annulus with one component in each of $R_\pm$; or*
2. *a disk which intersects γ in two components.*

Then $(M, \gamma, \beta) \overset{S}{\rightsquigarrow} (M', \gamma', \beta')$ is taut.

Proof. M' is evidently irreducible. In either case each of $R'_\pm$ has the same norm as $R_\pm$ so if it were not norm-minimizing one could find a smaller norm representative of $R_\pm$ contrary to the assumption that M is taut. ☐

4.5. **Windows and Guts.** In this section we sketch the proof of Theorem 4.18. The full argument involves a rather tedious analysis of a number of special cases, and it does not seem to make sense to reproduce it here when [4] and [12] are available.

A sutured manifold (M, γ, β) admits a decomposition into two pieces: a *window* W, which is an I bundle whose fibers run between R_- and R_+, and the *guts* G which is the remainder. This is closely analogous to the JSJ decomposition for a manifold with boundary. If we think of M as an orbifold with mirrors along γ that meet at right angles along β, we can take a regular manifold cover $\hat{M}$, perform the JSJ decomposition equivariantly in $\hat{M}$, and push the decomposition down to M.

Every component A of $A(\gamma)$ that meets some component of β bounds a nontrivial component of W, since a parallel copy of A pushed into M but without corners is not parallel to A as a surface with corners. In particular, no component of G intersects β.

Every component of G is a pared hyperbolic manifold with respect to γ (so that it admits a hyperbolic structure for which the sutures are cusps and $R_\pm$ are totally geodesic). Ignoring the sutures, it might admit compressing disks, but the boundary of every compressing disk must intersect an even number of sutures, and this number must be greater than 2 (or else we could split off a product square$\times I$ into W); thus as a surface with corners, each compressing disk has Euler characteristic a negative integer.

For simplicity let's suppose $M = G$ consisting of a single component. A *decomposing system* is a maximal set of non-parallel compressing disks $D = \{D_i\}$ meeting γ efficiently in such a way that every component of $G - D$ is either irreducible with incompressible boundary, or it is a B^3 meeting D in three components.

We shall define a complexity function associated to G consisting of an ordered pair of complexity functions. The first is the *height* $h(G)$ of the non-ball components of $G-D$; this is the maximum length of a partial hierarchy for $G-D$ using only non-disk surfaces that are incompressible and boundary incompressible in $G - D$ (see Chapter 1, Proposition 4.10). The second is the *index* $X(G) := \sum \chi(D_i)^2$ (compare Scharlemann [12], Definition 4.10); we may assume without loss of generality that our disk system D minimizes X. Now let's consider the complexity function for G consisting of the pair $(h(G), X(G))$ with the lexicographic ordering.

Lemma 4.23. *The set of values of the complexity function $(h(G), X(G))$ is well-ordered, and equals the minimum value $(0, 0)$ if and only if G is empty.*

Proof. Each of $h(G), X(G)$ is a non-negative integer, so the well-ordered property is immediate. If $h(G) = 0$ then G is a handlebody. If $X(G) = 0$ it admits a system of disks each meeting the sutures in two components; but this implies that G is a product sutured manifold, so that it is not contained in the guts after all. □

We shall show if $(h(G), X(G)) > (0, 0)$ that there is a special decomposing surface inducing a taut decomposition and strictly reducing the complexity.

Definition 4.24. Let E be an innermost disk of $D - S$ which becomes a proper disk E' in the sutured manifold M' obtained from M be decomposition along S. We say E is a *good disk* for S if $\chi(E') \ge 0$.

Note that $\chi(E')$ (and therefore whether E is good or not) depends not only on the topology of M, D, S but the orientation on S and $R_\pm$, since these determine the sutures of M', and therefore the corners of E'. Seven examples of good disks are illustrated in Figure 4.2 from [4].

Lemma 4.25. *Let S be a special decomposing surface for $M = G$ inducing a taut decomposition, and suppose that S cuts off a good disk $E \subset D$. Let S_E be the surface obtained by boundary compressing S along E. Then S_E induces a taut decomposition.*

Proof. If $\chi(E') > 0$ then one of $R'_\pm$ is compressible in M', contrary to the hypothesis that S is taut. Therefore $\chi(E') = 0$, and E' itself induces a taut decomposition of M' by Lemma 4.22. But this is precisely the result of decomposition along S_E. □

Proposition 4.26. *Let (M, γ, β) be a taut sutured manifold. Then either M is a product sutured manifold, or there is a special decomposing surface S inducing a taut decomposition that strictly decreases the complexity of the guts.*

Proof. If the intersection of S with the non-handlebody part of G is essential, decomposition along S reduces $h(G)$. Otherwise it induces a nontrivial decomposition of the handlebody part. The components of $D - S$ give rise to a decomposing system D' for G'. If there were a good disk E in $D - S$ we could have modified S to S_E by Lemma 4.25, thus without loss of generality we may assume there is no good disk. A disk D_i of D is decomposed into k subdisks $D'_{i,j}$ in D' where each $\|D'_{i,j}\| \geq 1$ and $\sum \|D'_{i,j}\| = \|D_i\|$ so that $X(G') < X(G)$ as claimed.

One must check that the part of S in the window does not induce a decomposition that creates new gut pieces, and that if the guts are nonempty, one may always find a special decomposing surface that intersects the guts essentially, but these facts are routine. □

This concludes our sketch of the proof of Theorem 4.18.

4.6. **Taut foliations from hierarchies.**

Definition 4.27. Let (M, γ, β) be a sutured manifold. A foliation $\mathcal{F}$ of M is *taut* if

(1) $\mathcal{F}$ is tangent to $R_\pm$ and transverse to γ;
(2) $\mathcal{F}|A(\gamma)$ is transverse to β and the I fibers;
(3) there is a proper compact 1-manifold X with $\partial X \subset R$ transverse to $\mathcal{F}$ and intersecting every leaf.

The main theorem of this section is

Theorem 4.28 (Gabai; taut foliation exists)**.** *A sutured manifold admits a taut foliation if and only if it is taut.*

Proof. One direction is easy. Let $\mathcal{F}$ be a taut foliation of (M, γ, β). Thinking of M as an orbifold with mirrors along γ and corners along β, we may take a manifold cover $\hat{M}$ and pull back $\mathcal{F}$ to a foliation $\hat{\mathcal{F}}$ on $\hat{M}$ transverse to a proper compact 1-manifold $\hat{X}$. Doubling $\hat{M}$ along its boundary gives rise to a closed manifold $D\hat{M}$ with a foliation $D\hat{\mathcal{F}}$ that is transverse to a proper compact 1-manifold $D\hat{X}$ and is therefore taut.

The preimages $\hat{R}_\pm$ become compact leaves of $D\hat{\mathcal{F}}$ and are therefore norm-minimizing in $D\hat{M}$ by Corollary 4.9; consequently $R_\pm$ are norm-minimizing in M. Furthermore, since $D\hat{M}$ is irreducible, the same is true of M. Thus (M, γ, β) is taut.

Now we show the converse. By Theorem 4.18, there is a taut sutured hierarchy for (M, γ, β) along special decomposing surfaces. A product sutured manifold obviously admits a taut (product) foliation. Thus to prove the theorem it suffices to show that if S decomposes (M, γ, β) to (M', γ', β') and (M', γ', β') admits a taut foliation $\mathcal{F}'$, then (M, γ, β) admits a taut foliation $\mathcal{F}$.

By Lemma 4.20 we may assume S is a special decomposing surface. If S is connected and $\partial S \subset \gamma$ then M may be obtained from M' by identifying $S'_\pm$ which are components of $R'_\pm$ and therefore leaves of $\mathcal{F}'$; The resulting foliation $\mathcal{F}$ is evidently taut if $\mathcal{F}'$ is. Thus we may assume that ∂S intersects $R_\pm$ nontrivially, and for each planar (resp. nonplanar) component the intersection is a family of parallel oriented homologically essential arcs (resp. loops).

The result of splitting along S produces new sutures; define $Z := \gamma' - \gamma = N(\partial S'_+ \cap R_-) \cup N(\partial S'_- \cap R_+)$. Each circle of $\partial S \cap R_\pm$ in a non-planar component of R gives rise to a new annulus in Z, and each proper arc of $\partial S \cap R_\pm$ in a planar component of R gives rise to a new rectangle in Z. See Figure 5.

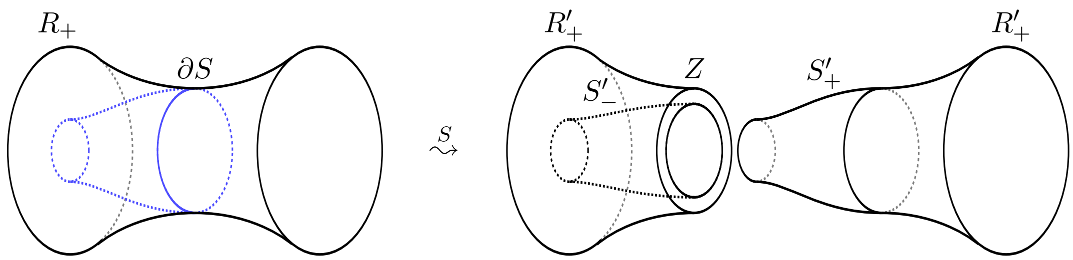

FIGURE 5. The surface S splits a non-planar component of R_+ in M', giving rise to a new annular suture $Z \subset \gamma'$.

The first step is to glue $S'_+ \subset R'_+$ to $S'_- \subset R'_-$. This produces a manifold M'' homeomorphic to M with a foliation $\mathcal{F}''$ that is transverse to $\partial M''$ along the image of γ'. See Figure 6. We must add a suitably foliated region to M'' to obtain M in such a way that $\mathcal{F}''|Z$ closes up in M to produce $\mathcal{F}$.

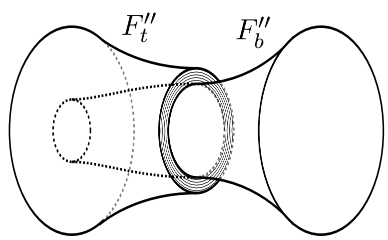

FIGURE 6. After gluing S'_+ to S'_- the part of the foliation $\mathcal{F}''$ transverse to Z is exposed.

It is clear that this modification can be performed independently on each component of $R_\pm$ intersecting ∂S. Let F be such a component, let F' be the image of F in M', and let F'' be F' in M''. We let $\partial_h Z$ be the part of ∂Z in $\partial F''$. If Z is an annulus (associated to a circle of $\partial S \cap R$ then $\partial_h Z = \partial Z$ consists of two circle components of $\partial F''$; otherwise $\partial_h Z$ consists of two arc components of $\partial F''$ ending at corners.

Because S is a special decomposing surface, $\partial S \cap F$ is a union of parallel, homologically essential circles or arcs in F. For simplicity we shall assume in the sequel that $\partial S \cap F$ consists of a single component.

Case 1: F'' is disconnected.
Let the two components of F'' be F''_t and F''_b where the subscripts stand for top and bottom, where the meaning is as indicated in Figure 6. Let $\partial_Z F''_b = \partial F''_b \cap \partial_h Z$; this is either an arc (ending at corners) or a circle. Since $\partial S \cap F$ is homologically essential but nevertheless separating, it pairs nontrivially with some proper arc with endpoints on both sides of ∂S; it follows that $\partial F''_b - \partial_Z F''_b$ is nonempty.

We attach $F_b'' \times I$ to M'' by identifying $F_b'' \times 0$ with F_b'' and $\partial_Z F_b'' \times I$ with Z, and let M be the result. It remains to extend the foliation $\mathcal{F}''|Z$ over $F_b'' \times I$ so that it is transverse to the I fibers and tangent to $F_b'' \times 1$.

Case 1a: $\partial_Z F_b''$ is an arc.
Since $\partial_Z F_b''$ is simply-connected, the foliation $\mathcal{F}''|Z$ is a product and we may extend $\mathcal{F}''|Z$ to a product foliation of $F_b'' \times I$.

Case 1b: $\partial_Z F_b''$ is a circle.
We construct a foliation of $F_b'' \times I$ from a representation $\rho : \pi_1(F_b'') \to \text{Homeo}^+(I)$. Since $\partial_Z F_b''$ is a circle and $\partial F_b'' - \partial_Z F_b''$ is nonempty it follows that F_b'' is a compact orientable surface with at least two boundary components, and therefore its fundamental group is free, and we may take $\partial_Z F_b''$ to be a free generator. We may therefore construct ρ arbitrarily on the other free generators (for instance, we can take it to be the identity element). This concludes the proof in Case 1.

Case 2: F'' is connected.
In this case $\partial_Z F''$ consists of two components, that we denote $\partial_b F''$ and $\partial_t F''$. Define $M_0 := M''$ and inductively we build M_n from M_{n-1} by attaching $F'' \times I$ to M_{n-1} by gluing $F'' \times 0$ to F'' and gluing $\partial_b F'' \times I$ to Z, and then relabeling so that $F'' \times 1$ becomes the new copy of F'' in M_n and $\partial_t F'' \times I$ becomes the new copy of Z in M_n.

We need to extend the foliation inductively over each $F'' \times I$ given the restriction to $\partial_b F'' \times I$. As before there are two cases (that $\partial_b F''$ is an arc, or that it is a circle) and as in Case 1 the extension is straightforward: a foliated bundle over an arc is trivial and may be extended as a trivial product; π_1 of a surface with (at least) two circle boundary components is free, and any of the circle boundary components may be taken to be a free generator. See Figure 7 for an example of a (trivially) foliated $F'' \times I$ where F is a pair of pants decomposed by an arc of ∂S.

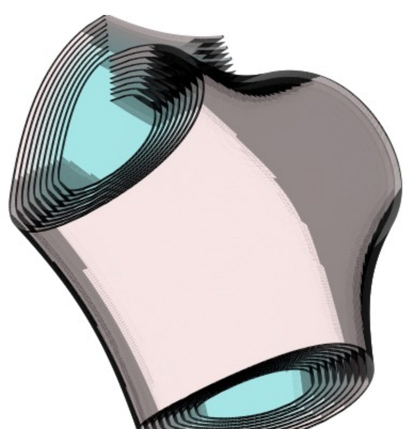

FIGURE 7. A foliated $F'' \times I$ where F is a pair of pants decomposed by an arc of ∂S.

This gives the extension over M_n for each n. Now take M_∞ to be the infinite union; this is an open manifold that may be compactified by adding a copy of F in such a way that the foliation spirals around F and extends to the compactification with F as a closed leaf. This compactification is M. This concludes the proof in Case 2.

This completes the induction step, and therefore the proof of the theorem. □

4.7. **Smooth and Finite Depth.** In the proof of Theorem 4.28 there is a great deal of choice in the choice of holonomy in the foliated I-bundles extending the foliation up the sutured manifold hierarchy at the induction step. A judicious sequence of choices allows one to construct taut foliations with certain additional properties.

Definition 4.29. Let $\mathcal{F}$ be a foliation. A leaf λ has *depth* 0 if it is compact, and has *depth* n if $\overline{\lambda} - \lambda$ is a union of leaves of depth $< n$ but not $< n-1$. A leaf is *finite depth* if it has depth n for some n.

A foliation has *depth* n if every leaf has depth $\le n$ and some leaf has depth n. A foliation has *finite depth* if it has depth n for some n.

Theorem 4.30 (Finite depth foliation)**.** *Let M be a compact irreducible, orientable 3-manifold, possibly with incompressible boundary, and let S be a Thurston-norm minimizing surface. Then there is a finite depth taut co-orientable foliation $\mathcal{F}$ with S as a compact leaf.*

Proof. One may construct a taut sutured hierarchy that begins by decomposing along S, and then build a taut foliation $\mathcal{F}$ as in the proof of Theorem 4.28. We show by induction that $\mathcal{F}$ can be chosen to be finite depth.

This is obvious at the base step, since a product foliation is depth 0. Let's suppose by induction, and with notation as in the proof of Theorem 4.28, that $\mathcal{F}'$ is a finite depth foliation of M'. In Case 1 the foliation $\mathcal{F}$ is obtained from $\mathcal{F}'$ by adding a foliated I bundle over a compact surface; in either case if we extend the holonomy by the identity on additional free factors, a leaf of $\mathcal{F}$ restricted to M' is equal to a leaf of $\mathcal{F}'$. Thus $\mathcal{F}$ is finite depth if $\mathcal{F}'$ is.

In Case 2 $\mathcal{F}$ is obtained from $\mathcal{F}'$ by adding countably many foliated I bundles over compact surfaces which spiral around a new closed leaf of $\mathcal{F}$. If each foliated I bundle has holonomy extended by the identity on additional free factors, then as before $\mathcal{F}$ is finite depth if $\mathcal{F}'$ is. □

A finite depth foliation is typically not smooth: in Case 2 where $\partial_b F''$ is a circle, and $\mathcal{F}'$ has nontrivial holonomy around this circle, and we extend the holonomy repeatedly by the identity on additional free factors, Kopell's Lemma (i.e. Theorem 1.9) says that $\mathcal{F}$ cannot be made C^2 near the closed boundary leaf.

When the genus of the boundary leaf is > 1, or if the boundary leaf itself has boundary, one may modify the construction as follows. We have a surface F'' and two distinguished boundary components $\partial_b F''$ and $\partial_t F''$. Let $\alpha \in \text{Diff}^+(I)$ be the holonomy of the foliated I bundle around $\partial_b F''$ in the foliated annulus $Z \subset M_0$. If F'' has genus g with h boundary components, we may choose free generators $a_1, b_1, \cdots, a_{g-1}, b_{g-1}, c_1, \cdots, c_{h-1}$ for $\pi_1(F'')$ where c_1 represents the loop $\partial_b F''$ and $\prod_{j=1}^{h-1} c_j \prod_{j=1}^{g-1} [a_j, b_j]$ represents the loop $\partial_t F''$.

If $h > 2$ (equivalently, if ∂F is nonempty) then by defining ρ suitably on c_{h-1} we may ensure that the foliation is trivial on $\partial_t F''$, and then this foliation may be extended as a product over each successive $F'' \times I$ factor and spun (smoothly) around the limit leaf F;

in fact, if we choose suitably coordinates near F we may arrange that the foliation is C^∞ tangent to the identity along F.

If $h = 2$ but $g > 1$ then the holonomy around $\partial_b F''$ and $\partial_t F''$ differ by a product of $g-1$ commutators. The group $\text{Diff}^+(I)$ is not perfect; however the subgroup Γ of diffeomorphisms C^∞ tangent to the identity at 0 and 1 is:

Theorem 4.31 (Sergeraert)**.** *Let Γ be the group of orientation-preserving diffeomorphisms of I that are C^∞ tangent to the identity at the endpoints. Then Γ is perfect.*

Proof. □

See [13] for a proof.

From this we may conclude:

Theorem 4.32 (Smooth foliation)**.** *Let M be a compact irreducible, atoroidal orientable 3-manifold, possibly with incompressible boundary, and let S be a Thurston-norm minimizing surface which either has boundary, or is closed of genus at least 2. Then there is a C^∞ taut co-orientable foliation $\mathcal{F}$ with S as a compact leaf.*

We sketch the proof.

Proof. Fix notation as in the proof of Theorem 4.28. Suppose by induction that $\mathcal{F}'$ is C^∞ and the germ of the holonomy around every compact leaf is C^∞ tangent to the identity. When we glue up M' to produce M'' the holonomy pieces together smoothly along F', and α as above is C^∞ tangent to the identity at 0 and 1. Thus we may write α as a product of n commutators in Γ for some n. Let $m = \lceil n/(g-1) \rceil$. Observe that $M_m - M_0$ consists of m copies of $F'' \times I$ concatenated end to end; i.e. it is of the form $\Sigma \times I$ where Σ has genus $m(g-1)$ (and possibly multiple boundary components if $h > 2$). It follows from Theorem 4.31 that we can choose a smooth foliation on $M_m - M_0$ that is C^∞ tangent to the identity along the boundary leaves, which fits together with $\mathcal{F}''$ along Z in M_0, and which is foliated as the identity along the new copy of Z in M_m. This may be extended as a product foliation over M_∞ spiraling around a compact leaf F. If we choose suitable coordinates on the end of M_∞ the germ of the holonomy of the resulting foliation is C^∞ tangent to the identity along F. This completes the induction step and proves the theorem. □

4.8. **Branched surfaces.** An equivalent way to think of a sutured manifold is to collapse the annular regions to (oriented) polygonal curves, collapsing β in this way to vertices, and to comb the boundary into a branched surface with corners, branched along the sutures. See Figure 8. In this picture, if S is a properly embedded decomposing surface S in (M, γ), the boundary ∂S gets corners where it is tangent to $\partial R_\pm$ and cusps where it is transverse to $\partial R_\pm$.

5. Holomorphic geometry

There are two key theorems in the classical theory of Riemann surfaces. The first, the Uniformization Theorem of Koebe–Poincaré–Klein, says that any Riemann surface admits a metric of constant curvature in its conformal class, unique up to scaling. The second is Poincaré's theorem that every Riemann surface is algebraic, equivalently it may be holomorphically embedded in $\mathbb{CP}^n$ for some n.

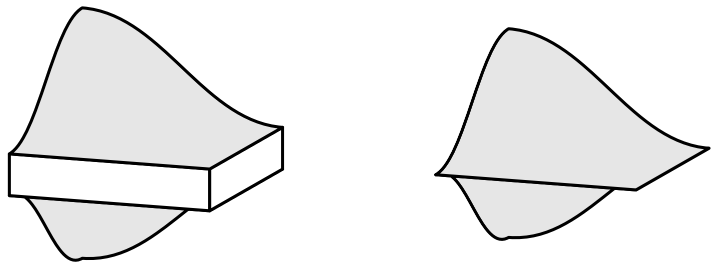

FIGURE 8. The boundary of a sutured manifold collapses along $A(\gamma)$ to a branched surface.

Somewhat remarkably both of these theorems have their analogs in the theory of taut foliations of 3-manifolds.

5.1. **Uniformization.** Suppose $\mathcal{F}$ is an oriented, co-oriented foliation of a 3-manifold M. A Riemannian metric on M restricts to a Riemannian metric on each leaf, giving it an associated holomorphic structure. An elliptic leaf is topologically S^2, so if $\mathcal{F}$ contains such a leaf by the Reeb Stability Theorem 1.21 M is $S^2 \times S^1$ and $\mathcal{F}$ is the product foliation.

If $\mathcal{F}$ taut contains a parabolic leaf then necessarily M is toroidal:

Proposition 5.1 (Parabolic implies toroidal)**.** *If $\mathcal{F}$ taut contains a parabolic leaf then M contains a homologically essential incompressible torus.*

Proof. Since $\mathcal{F}$ is taut, M is irreducible. Let λ be a parabolic leaf of $\tilde{\mathcal{F}}$ covering a parabolic leaf of $\mathcal{F}$. We have already seen in the proof of Proposition 3.14 that there is a sequence of subdisks λ_i with

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

so that $\lambda_i/\text{area}\lambda_i$ converges to the homology class $[\mu] \in H_2(M;\mathbb{R})$ of a nontrivial invariant transverse measure μ.

Since M is compact, there is a constant C so that any homologically trivial loop γ in M bounds an immersed surface S for which $\text{area}(S)$, $\text{diameter}(S)$ and $|\chi(S)|$ are bounded by $C\,\text{length}(\gamma)$.

Let $\pi(\lambda_i)$ be the projection of λ_i to M, and let S_i be an immersed surface bounding $\pi(\partial\lambda_i)$ and satisfying the inequality above. Then $F_i := S_i \cup \pi(\lambda_i)$ is a sequence of closed immersed surfaces with $[F_i]/\text{area}(\lambda_i) \to [\mu]$ and $|\chi(F_i)|/\text{area}(\lambda_i) \to 0$. It follows that the Thurston norm of the class $[\mu]$ is zero, and therefore that some homology class in M is toroidal. □

Consequently, if M is atoroidal, any taut foliation $\mathcal{F}$ of M has all leaves conformally hyperbolic.

Theorem 5.2 (Candel, Uniformization)**.** *Let $\mathcal{F}$ be a foliation of M. Then M has a C^0 Riemannian metric in any conformal class for which every leaf of $\mathcal{F}$ has constant curvature -1 if and only if every invariant transverse measure μ has $\chi([\mu]) < 0$.*

Proof. A measure supported on spherical leaves obviously has $\chi > 0$ and we have just seen that a parabolic leaf gives rise to a measure with $\chi = 0$. If $\mathcal{F}$ is a foliation with all leaves hyperbolic and μ is an invariant transverse measure, the Gauss–Bonnet formula gives $\chi < 0$. So we are left to show that if every leaf of $\mathcal{F}$ is conformally hyperbolic, then M has a metric for which every leaf of $\mathcal{F}$ is individually hyperbolic.

Fix a Riemannian metric g on M. Since each leaf is conformally hyperbolic, there is a unique function $f \to \mathbb{R}^+$ which is leafwise smooth and so that fg is leafwise hyperbolic. Our task is to show that f is continuous.

Let's work in the universal cover for simplicity. Fix a point $p \in \lambda$ and let $p_i \in \lambda_i$ converge to p. Let $\mathbb{D}$ denote the open unit disk in $\mathbb{C}$. For each i there is a conformal isomorphism $\phi_i : \mathbb{D} \to \lambda_i$ taking 0 to p_i, unique up to precomposing with a rotation. Give $\mathbb{D}$ its hyperbolic metric, so that $f(p_i) = |d\phi_i(0)|^{-2}$.

If the maps ϕ_i were equicontinuous we could extract a limit of any subsequence. This can only fail if the derivatives of the ϕ_i blow up somewhere. If the derivative of ϕ_i blows up near infinity we use a trick due to Brody [1]: we may precompose ϕ_i with $\psi_\epsilon : \mathbb{D} \to \mathbb{D}$ which takes p to $(1-\epsilon)p$ for some small ϵ so that the derivative of $\phi_i\psi_\epsilon$ is maximized at some point in the interior. After precomposition with another Möbius transformation, we may assume $|d(\phi_i\psi_\epsilon)|$ is maximized at 0. If this maximum still blew up as $i \to \infty$, we could extract in the limit (by rescaling the domain so that the derivative at 0 has norm 1) a nonconstant conformal map $\phi : \mathbb{C} \to \mu$ with image contained in some leaf μ of $\mathcal{F}$. But this is impossible if μ is conformally hyperbolic, and therefore the maps $\phi_i\psi_\epsilon$ are equicontinuous after all for any fixed $\epsilon > 0$. Thus after passing to a diagonal subsequence (taking $\epsilon \to 0$) we deduce that the ϕ_i converge on compact subsets to a nonconstant conformal map $\phi : \mathbb{D} \to \lambda$. The map ϕ is not necessarily a covering map (or even surjective) so by the Schwarz Lemma the conformal covering map $\phi_\lambda : \mathbb{D} \to \lambda$ has a bigger derivative at 0 than ϕ. Said another way, $f(p) \le \liminf f(p_i)$ so that f is lower semicontinuous.

Conversely, if $p \in \lambda$ and $\phi_\lambda : \mathbb{D} \to \lambda$ is the uniformizing map, for any ϵ the map $\phi_\lambda\psi_\epsilon$ may be approximated by conformal maps from $\mathbb{D}$ to λ_i taking 0 to p_i. It follows that f is upper semicontinuous, and therefore continuous. It follows a posteriori from the Schwarz Lemma that the uniformizing maps $\phi_i : \mathbb{D} \to \lambda_i$ as above actually converge to $\phi : \mathbb{D} \to \lambda$. □

Remark 5.3. It is not necessarily true that the metric provided by Theorem 5.2 is smooth in transverse directions (though one can show it is Hölder continuous for some exponent). However it is true that the partial derivatives of all orders in leafwise directions are continuous on M. This follows from the Cauchy integral formula given what we have already proved.

By mollifying the metric one can arrange for g to be smooth on M and have leafwise curvature pinched between $-1 \pm \epsilon$.

5.2. **Projective Embeddings.**

Theorem 5.4 (Ghys, Ample)**.** *A smooth co-orientable foliation $\mathcal{F}$ of M is taut if and only if there is some integer n, and an embedding $\varphi : M \to \mathbb{CP}^n$ which is holomorphic 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 restrict to local 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. First suppose for simplicity that every leaf is conformally hyperbolic, so that by Candel's Theorem 5.2 we may assume that each leaf is isometric to $\mathbb{H}^2$. Let X be a smooth section of $UT_\gamma\mathcal{F}$. In the unit disk model for $\mathbb{H}^2$, let v be a unit vector at the origin, and let $R(z)dz^2$ be a quadratic holomorphic differential where $R(z) = P(z)/z^k$ for some polynomial $P(z)$ of degree at most k.

If we lift to the universal cover, for each point p in a lift of γ contained in some leaf λ of $\tilde{\mathcal{F}}$, there is a unique isometry $\lambda \to \mathbb{H}^2$ taking p to 0 and $\tilde{X}(p)$ to v. Pulling back defines a quadratic holomorphic differential on λ. Summing leafwise over all transverse lifts of γ gives a global holomorphic section of $L^{\otimes k}$. Taking k big enough gives many sections, evidently enough to separate points of M.

If some leaf is conformally spherical, by Reeb stability M is $S^2 \times S^1$ foliated by spheres, and we may take φ to be projection to $S^2 = \mathbb{CP}^1$. If there is a mixture of hyperbolic and parabolic leaves we take a regular branched cover of M over γ and pull back $\mathcal{F}$ to obtain a foliation $\mathcal{G}$ of a new 3-manifold N whose leaves are now all hyperbolic. We may construct leafwise quadratic holomorphic differentials on $\mathcal{G}$ as above, and average them under the deck group of the cover so that they descend to differentials on $\mathcal{F}$. □

6. Universal circles

The universal cover of a hyperbolic surface is conformally equivalent to the unit disk, and may be canonically compactified by added a circle at infinity. The action of the deck group extends continuously to an action on this circle.

If $\mathcal{F}$ is a foliation of M without spherical or parabolic leaves (for instance if M is atoroidal and $\mathcal{F}$ is taut) Candel's Uniformization Theorem 5.2 says that M admits a (C^0) Riemannian metric for which every leaf becomes simultaneously a hyperbolic surface. If $\tilde{\mathcal{F}}$ is the pulled back foliation of the universal cover $\tilde{M}$ then every leaf λ of $\tilde{\mathcal{F}}$ may be compactified with a circle at infinity $S^1_\infty(\lambda)$ and the deck group of M acts on this collection of circles.

A *universal circle* is a single 'master circle' S^1_{univ} that collates and unifies the distinct $S^1_\infty(\lambda)$ in a coherent way so that the deck group (i.e. $\pi_1(M)$ becomes a group of automorphisms of S^1_{univ}. Before giving a precise definition we must introduce some additional terminology; this is done in the next two subsections.

6.1. Circle bundle E.

Let $\mathcal{F}$ be a taut foliation of M, and let L denote the leaf space of $\tilde{\mathcal{F}}$. Novikov's Theorem 2.4 implies that L is a connected, simply-connected 1-manifold; the catch is that L is typically not Hausdorff.

We say that two distinct leaves $\mu, \lambda \in L$ are *comparable* if there is a transversal τ to $\tilde{\mathcal{F}}$ from μ to λ; equivalently if there is an embedding of an interval $I \to L$ whose endpoints are taken to μ and λ respectively.

If $\mathcal{F}$ is co-orientable, then a choice of co-orientation determines an orientation on L and (since L is simply connected) a partial order, defined on comparable pairs of leaves, where $\mu < \lambda$ if and only if there is a positively oriented transversal τ from μ to λ. The action of $\pi_1(M)$ on L preserves the orientation, and therefore the partial order and the notion of comparability.

Now suppose that the leaves of $\mathcal{F}$ are all conformally hyperbolic, so that we may simultaneously uniformize them with a suitable metric on M. Each $\lambda \in L$ determines a circle $S^1_\infty(\lambda)$ and we let E denote the union of these circles.

One may topologize E as follows. A transversal τ to $\tilde{\mathcal{F}}$ projects to an interval in L, and we may denote the associated family of circles as $E|\tau$. On the other hand, the unit tangent bundle $UT\mathcal{F}$ restricts along τ to a circle bundle $UT\mathcal{F}|\tau$ with total space homeomorphic to a cylinder. For each leaf λ intersecting τ the restriction $UT\mathcal{F}|\tau \cap \lambda$ is a circle. Each vector in this circle is tangent to a unique oriented geodesic ray in λ, that limits to a unique point in $S^1_\infty(\lambda)$; thus we obtain an 'endpoint map' $e : UT\mathcal{F}|\tau \cap \lambda \to S^1_\infty(\lambda)$ that is evidently a homeomorphism. The family of maps as we vary over τ is a bijection $e : UT\mathcal{F}|\tau \to E|\tau$ and we may topologize $E|\tau$ by declaring that this map is a homeomorphism.

Since leaves vary continuously on compact subsets, this homeomorphism is independent of the choice of transversal τ; varying over all transversals defines a topology on E for which it becomes a circle bundle over L in the usual sense.

6.2. **Monotone maps.** The universal circle does not relate to the individual circles $S^1_\infty(\lambda)$ by homeomorphisms, but by a slightly weaker relation, that of a *monotone map*:

Definition 6.1 (Monotone map)**.** A map $f : S^1 \to S^1$ is *monotone* if f has degree 1 and the point preimages of f are contractible; i.e. for all $p \in S^1$ either $f^{-1}(p)$ is a single point, or $f^{-1}(p)$ is a connected interval.

If $f : S^1 \to S^1$ is monotone, it lifts to $\tilde{f} : \mathbb{R} \to \mathbb{R}$ where the monotone property says that if $p, q \in \mathbb{R}$ and $p \geq q$ then $\tilde{f}(p) \geq \tilde{f}(q)$.

Definition 6.2 (Core)**.** If $f : S^1 \to S^1$ is monotone the *core* of f, denoted $\text{core}(f)$, is the set of points where f is not locally constant.

The complement of $\text{core}(f)$ consists of countably many open intervals called the *gaps*. The value of f is constant on the closure of each gap.

Lemma 6.3 (Perfect)**.** *For any monotone $f : S^1 \to S^1$ the set* $\text{core}(f)$ *is perfect; i.e. it has no isolated points.*

Proof. If $x \in \text{core}(f)$ is isolated, it is in the closure of two distinct gaps, I and J say. But then $f|I = f(x) = f|J$ so that f is locally constant at x contrary to definition. □

Lemma 6.4 (Semicontinuity)**.** *The core of a family of monotone maps varies lower semicontinuously. That is, if $f_i : S^1 \to S^1$ converges in the compact-open topology to $f : S^1 \to S^1$ and $x \in \text{core}(f)$ there are $x_i \in \text{core}(f_i)$ with $x_i \to x$.*

Proof. Since $x \in \text{core}(f)$, for any open interval I containing x the values of f on $\partial\bar{I}$ are distinct. If $f_i \to f$ then the values of f_i on $\partial\bar{I}$ are likewise eventually distinct for sufficiently large i, thus there is some $x_i \in I \cap \text{core}(f_i)$. □

If F is a path-connected family of monotone maps, we let $\text{core}(F)$ denote the closure of $\cup_{f\in F}\text{core}(f)$. Note that $\text{core}(F)$ is perfect by Lemma 6.3. Two closed subsets X and Y of S^1 are said to be *unlinked* if there do not exist disjoint two element subsets $A \subset X$ and

$B \subset Y$ that are linked as a pair of S^0s in S^1. Note that the perfectness of cores implies that core(F) and core(G) are unlinked if and only if core(F) is contained in the closure of a single gap of core(G) and vice versa.

Proposition 6.5 (Unlinked families)**.** *Let F and G be two path-connected families of monotone maps. Suppose for all $f \in F$ and $g \in G$ that* core(f) *and* core(g) *are unlinked. Then* core(F) *and* core(G) *are unlinked.*

Proof. We first show that if we fix $g \in G$ that core(F) and core(g) are unlinked. There are countably many gaps I_i for core(g), and by hypothesis for all $f \in F$ there is an index i so that core$(f) \subset \bar{I}_i$. Since core(g) is perfect, the $\bar{I}_i$ are disjoint, and therefore the subset F_i of F with core$(f) \subset \bar{I}_i$ is well-defined.

We claim that each F_i is closed. For, if $f_j \in F_i$ so that there are $x_j \in$ core$(f_j) \cap \bar{I}_i$ and $f_j \to f$ but $x \in$ core$(f) \cap \bar{I}_k$ for $i \neq k$ then by Lemma 6.4 there are $y_j \in$ core(f_j) converging to x. But then y_j is not in $\bar{I}_i$ for sufficiently large j contrary to the fact that core(f_j) and core(g) are unlinked. This contradiction proves the claim.

A Theorem of Sierpinski [14] says that a path-connected set does not admit a nontrivial decomposition into countably many closed subsets; thus $F = F_i$ for some i and core(F) and core(g) are unlinked.

The same argument, replacing g with F and F with G, shows that core(F) and core(G) are unlinked. □

6.3. **Universal Circles.** We are now in a position to give the definition of a universal circle. For simplicity we give the definition for $\mathcal{F}$ co-oriented.

Definition 6.6. Let $\mathcal{F}$ be a co-oriented foliation of M with conformally hyperbolic leaves, and let L denote the leaf space of $\widetilde{\mathcal{F}}$. A *universal circle* for $\mathcal{F}$ consists of

(1) an oriented circle S^1_{univ} and a representation $\phi_{\text{univ}} : \pi_1(M) \to \text{Homeo}^+(S^1)$;
(2) for each $\lambda \in L$ a monotone map $\pi_\lambda : S^1_{\text{univ}} \to S^1_\infty(\lambda)$ depending continuously on λ

such that

(1) for all $\lambda \in L$ and $\alpha \in \pi_1(M)$ there is a commutative diagram:

$$\begin{array}{ccc} S^1_{\text{univ}} & \xrightarrow{\phi_{\text{univ}}(\alpha)} & S^1_{\text{univ}} \\ \downarrow \pi_\lambda & & \downarrow \pi_{\alpha(\lambda)} \\ S^1_\infty(\lambda) & \xrightarrow{\alpha} & S^1_\infty(\alpha(\lambda)) \end{array}$$

(2) if μ and λ are incomparable then core(π_μ) and core(π_λ) are unlinked in S^1_{univ}.

The next few subsections will be devoted to a proof of the following

Theorem 6.7 (Thurston, Calegari–Dunfield; Universal Circles)**.** *Let $\mathcal{F}$ be a co-oriented foliation of M with conformally hyperbolic leaves. Then $\mathcal{F}$ admits a universal circle.*

6.4. **Leaf Pocket Theorem.** Our first step is to compare the behavior near infinity of nearby leaves of $\widetilde{\mathcal{F}}$. It will turn out, for each leaf λ, that there is a dense set of $S^1_\infty(\lambda)$ so that geodesic rays in λ asymptotic to this subset stays 'close' to geodesic rays in all sufficiently nearby leaves.

This statement is complicated by the fact that there is typically no uniform comparison between the intrinsic and extrinsic geometry of leaves of $\tilde{\mathcal{F}}$ in $\tilde{M}$; we must therefore restrict attention to neighborhoods in which such comparisons can be made.

Definition 6.8 (Separation Constant)**.** A number $\epsilon > 0$ is a *separation constant* for $\tilde{\mathcal{F}}$ if there is some L so that for every leaf λ of $\tilde{\mathcal{F}}$ the inclusion map from λ into its ϵ-neighborhood $N_\epsilon(\lambda)$ is L-bilipschitz in the respective path metrics. In other words, for all λ, and all $p, q \in \lambda$ if there is a path in $N_\epsilon(\lambda)$ from p to q of length at most T then there is a path in λ from p to q of length at most LT.

Lemma 6.9. *Every foliation of a compact manifold admits a separation constant.*

Proof. In any product chart one may project locally to any leaf by a uniformly Lipschitz map. Since M may be covered by finitely many product charts, the lemma follows. □

Definition 6.10 (Marker)**.** Let λ_t for $t \in [0, 1]$ be an interval of leaves of $\tilde{\mathcal{F}}$. A *marker* for this family of leaves is a map $m : [0, 1] \times \mathbb{R}^+ \to \tilde{M}$ such that

(1) for each $p \in [0, 1]$ the map $m : p \times \mathbb{R}^+ \to \lambda_p$ is a (not necessarily parameterized) geodesic ray;
(2) for each $t \in \mathbb{R}^+$ the transversal $m([0, 1] \times t)$ has length $\leq \epsilon/2$.

The *endpoint* of m, denoted $e(m)$ or just e if m is understood, is the interval in E consisting of the endpoint in each $S^1_\infty(\lambda_p)$ of the ray $m(p \times \mathbb{R}^+)$.

Theorem 6.11 (Leaf Pocket Theorem)**.** *Let $\mathcal{F}$ be taut with hyperbolic leaves. There is a $\pi_1(M)$-equivariant collection $\mathcal{M}$ of markers so that*

(1) endpoints of distinct markers are either disjoint in E or their union is an embedded interval; and
(2) for all λ the intersection of the endpoints of markers with $S^1_\infty(\lambda)$ is dense in $S^1_\infty(\lambda)$.

Proof. As in the proof of Proposition 3.15, every minimal set Λ of $\mathcal{F}$ contains a non-simply connected leaf λ. Let $\gamma \subset \lambda$ be a simple nontrivial geodesic. If τ is a sufficiently short transversal with one endpoint on γ then for some choice of orientation, holonomy transport of τ around γ takes τ into itself. The transversal τ sweeps out a rectangle, that glues up to make a 'sawblade' S transverse to $\mathcal{F}$; see Figure 9.

Lift S to the universal cover and straighten it to a geodesic ray leafwise. If τ is sufficiently short, an interval of geodesic rays contained in the straightened sawblade is a marker; in other words every lift of S is a union of markers, whose endpoints in E piece together to make an embedded interval. Choose one sawblade for each minimal set, with τ short enough in each case that distinct sawblades do not intersect, and then take as $\mathcal{M}$ the union of the lifts.

We claim that distinct lifts of sawblades are never asymptotic at infinity in any leaf. Let λ be a leaf, and suppose $\tilde{S}, \tilde{S}'$ are two lifts of sawblades S and S' intersecting λ (we allow the possibility that $S = S'$). Let $\gamma = \tilde{S} \cap \lambda$ and $\gamma' = \tilde{S}' \cap \lambda$. These are geodesic rays in λ, disjoint because S and S' are either equal or disjoint. If γ and γ' had a common point at infinity, S and S' would accumulate on each other, contrary to the fact that they are compact and disjoint. This proves the claim, and the first bullet point of the theorem.

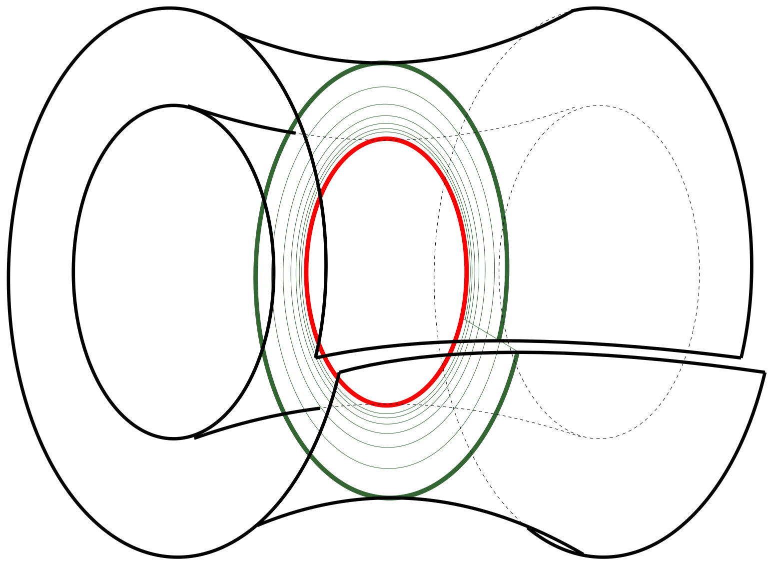

FIGURE 9. A simple closed geodesic (in red) on a leaf bounds a sawblade (in green) on either side.

Now we prove the second bullet point. Let's consider a fixed minimal set Λ and a sawblade S. Since Λ is minimal, S intersects every leaf of Λ. Since Λ is compact there is a T so that every point in every leaf of Λ may be joined by a path of length $\le T$ to a point in the interior of S. It follows that for every leaf λ of $\tilde{\Lambda}$ that the intersection with lifts of S is a union of geodesic rays that come within distance T of every point. We have already shown that no two of these rays have common endpoints at infinity; it follows that the set of endpoints of $\mathcal{M}$ in $S^1_\infty(\lambda)$ consists of at least two points. By compactness of Λ there is a least non-negative number θ strictly less than 2π so that for every λ in $\tilde{\Lambda}$, for every $p \in \lambda$ and for every interval I in $S^1_\infty(\lambda)$ with visual measure at least θ as seen from p, there is at least one endpoint of a sawblade. But *every* interval in $S^1_\infty(\lambda)$ has visual measure as close to 2π as we like as seen from some point in λ; it follows that $\theta = 0$, so that the endpoints of $\mathcal{M}$ are dense in $S^1_\infty(\lambda)$ for every λ in any $\tilde{\Lambda}$.

Now, if $\mu \in \tilde{\mathcal{F}}$ is arbitrary, and p_i is any sequence in μ limiting to some $p \in S^1_\infty(\mu)$, the balls $B_i(p_i)$ in μ of radius i project in $\mathcal{F}$ to a family of subsets containing a subsequence whose limit contains some minimal set Λ. It follows that infinitely many of the $B_i(p_i)$ intersect many lifts of a sawblade S for Λ, and therefore $S^1_\infty(\mu)$ intersects some endpoint of some marker in every neighborhood of p. Since μ and p are arbitrary, we are done. $\square$

6.5. **Leftmost sections.** Once we have defined the universal circle, each point $p \in S^1_{\text{univ}}$ will define a section σ_p of the circle bundle $E \to L$, by $\sigma_p(\lambda) = \pi_\lambda(p)$. Turning this around, we shall construct S^1_{univ} by first constructing (some of) these sections, and then taking a suitable completion. This construction will take several steps.

The first step is construct such sections over embedded intervals in L. Suppose we have constructed a family of markers $\mathcal{M}$ as in Theorem 6.11. The endpoints of $\mathcal{M}$ form a collection of disjoint intervals $\mathcal{E}$ in the cylinder $E|I$, transverse to the foliation by circles and intersecting every circle in a dense subset. Let's parameterize the interval I by $[0,1]$ so that the orientation on I agrees with the orientation on L, and let λ_t denote the leaves of I.

Definition 6.12 (Admissible section)**.** A section $\sigma : I \to E|I$ is *upwards (resp. downwards) admissible* if it satisfies the following two properties:

(1) if $\sigma(0) \in e$ for some $\mathcal{E}$ (resp. $\sigma(1) \in e$ for some $\mathcal{E}$) then $\sigma(I)$ contains $e \cap E|I$; and
(2) $\sigma(I)$ does not cross any element of $\mathcal{E}$.

We claim that among all upwards (resp. downwards) admissible sections with $\sigma(0) = p \in S^1_\infty(\lambda_0)$ (resp. $\sigma(1) = p \in S^1_\infty(\lambda_1)$), there is a unique *leftmost* section — one that moves positively (resp. negatively) around the circle as fast as possible as t increases from 0 to 1 (resp. decreases from 1 to 0).

7. Essential laminations

Definition 7.1. A *lamination* Λ in a 3-manifold M is a closed union of surfaces (leaves) so that M may be covered by coordinate charts (product charts) that intersect the leaves locally in horizontal disks.

Example 7.2. A foliation is a lamination. A closed surface is a lamination. A minimal set in a foliation is a lamination.

Let $\Lambda \subset M$ be a lamination. Define M_Λ to be the metric completion of $M - \Lambda$ with respect to the induced path-metric. Typically M_Λ is not compact, although it has boundary which maps to *boundary leaves* of Λ (those that are locally isolated on at least one side) under the obvious immersion $M_\Lambda \to M$.

Definition 7.3. A lamination is *essential* if no leaf is a sphere or a torus bounding a solid torus, and if M_Λ is irreducible, and admits no compressing disk or monogon.

Example 7.4. If M is irreducible, any essential surface in M is an essential lamination. If $\mathcal{F}$ is taut, any closed union of leaves of $\mathcal{F}$ is an essential lamination.

Example 7.5. Let M fiber over the circle with fiber F and monodromy ϕ. If ϕ is pseudo-Anosov the stable and unstable laminations of ϕ suspend to essential laminations of M.

The manifold M_Λ is typically non-compact; the ends correspond to regions in the interstices between Λ where two boundary leaves get very close. These ends may be covered with product charts in an obvious way giving them the structure of I-bundles over (non-compact) surfaces, that are partially compactified by finitely many *interstitial annuli*. The complement of the I-bundle regions comprises the *guts* of M_Λ, denoted $G(\Lambda)$. This is a compact manifold whose boundary decomposes into interstitial annuli, and compact subsurfaces of boundary leaves of Λ. If Λ is co-oriented, $G(\Lambda)$ is a sutured manifold.

7.1. **Laminar branched surfaces.** Recall from Chapter 1, § 6.2 the definition of a branched surface B, and what it means for B to carry a surface S. There is an obvious sense of what it means for B to carry a (nowhere dense) lamination Λ, i.e. we may isotop Λ so that the leaves of Λ run locally nearly parallel to B, as in Chapter 1 Figure 10. We say that B *fully carries* Λ if every sector of B is in the image of Λ.

Definition 7.6. A branched surface B is

(1) *essential* if it is incompressible (see Chapter 1, Definition 6.6.) and all complementary regions are irreducible;
(2) *taut* if it is co-orientable, if complementary regions are taut sutured manifolds, and through every sector there is a closed oriented curve positively transverse to B;
(3) *laminar* if it is essential, and has no trivial bubbles, sink disks or half sink disks.

Every taut branched surface is essential. Every lamination fully carried by an essential branched surface is essential.

Example 7.7 (Sutured manifold sequence to branched surface)*.* Let $\mathcal{S}$ be a sequence of sutured manifold decompositions

$$(M_0, \gamma_0) \overset{S_0}{\rightsquigarrow} (M_1, \gamma_1) \overset{S_1}{\rightsquigarrow} \cdots \overset{S_{n-1}}{\rightsquigarrow} (M_n, \gamma_n)$$

There is a branched surface $B(\mathcal{S})$ built from the union of the S_i, co-oriented compatibly with the sutured structures.

Theorem 7.8 (Li; laminar branched surfaces carry)**.** *A laminar branched surface fully carries an essential lamination.*

Corollary 7.9. *If B is taut and laminar, then B carries an essential lamination that is a sublamination of a taut co-orientable foliation.*

Example 7.10*.* Suppose M is irreducible with $H_2(M;\mathbb{Z})$ nontrivial, and let $\mathcal{S}$ be a taut sutured hierarchy for M. Then $B(\mathcal{S})$ is taut and laminar.

Example 7.11 (Dunfield; Foliar Orientations)*.* Let τ be a triangulation of M. An orientation on the edges is *acyclic* if the induced orientation on the edges of each simplex induces a total ordering of the vertices. For an acyclic orientation each simplex has a unique *longest edge* running from the smallest to the largest vertex (in the ordering). An acyclic orientation has *no sink edges* if there is no edge which is longest in every simplex it is contained in. An acyclic orientation is dual to a canonical co-oriented branched surface B, that intersects each simplex as in Figure **??**. The orientation is *taut* if each vertex is contained in a taut sutured ball of $M - B$.

Definition 7.12*.* An edge orientation is *foliar* if it is acyclic, taut, and admits no sink edges.

The no sink condition implies that B is laminar, and the taut condition implies that B is taut. Thus by Corollary 7.9 any 3-manifold that admits a foliar edge orientation contains a taut co-orientable foliation.

7.2. Persistently foliar knots.

Definition 7.13. Let M be a closed 3-manifold. Let $K \subset M$ be a knot, and let M_K denote the complement in M of an open solid torus neighborhood of K. A knot $K \subset M$ is said to be *persistently foliar* if for every non-meridian slope on ∂M_K there is a taut co-oriented foliation of M_K intersecting ∂M_K transversely in circles of the given slope.

It follows that if $K \subset M$ is persistently foliar, every manifold N obtained from M by non-trivial surgery along K admits a taut co-oriented foliation.

The purpose of this section is to prove the following theorem:

Theorem 7.14 (Delman–Roberts). *Every prime, non-torus alternating knot in S^3 is persistently foliar.*

Definition 7.15 (Double-diamond taut).

Proposition 7.16. *Let $\mathcal{S}$ be a sequence of taut sutured manifold decompositions*

$$(M, \partial M) \overset{R}{\rightsquigarrow} (M_1, \gamma_1) \overset{S}{\rightsquigarrow} (M_2, \gamma_2)$$

so that

(1) ∂M is a torus;
(2) ∂R is connected and nonempty; and
(3) S is double-diamond taut with respect to $\alpha\tau\alpha' \subset \partial S$.

Let σ be the essential simple loop on ∂M that is the union of τ and an arc of ∂R, and let M' be the 3-manifold obtained by Dehn filling M along σ, so that M is M' minus a neighborhood of a knot $K \subset M'$. If $B(\mathcal{S})$ has no sink disks disjoint from ∂M then K is persistently foliar.

Proof. If S is double-diamond taut, we show how to modify the branched surface so that the gut region in M' bounding K is a solid torus with two meridianal sutures. This will prove the Proposition, since for any nontrivial surgery on K this region becomes taut. □

8. RFRS and the Virtual Fibration Conjecture

9. Acknowledgments

Danny Calegari was supported by NSF grant DMS 1005246.

References

[1] R. Brody, *Compact Manifolds and Hyperbolicity*, Trans. AMS, **235** (1978) 213–219
[2] A. Candel, *Uniformization of surface laminations*, Ann. Sci. École Norm. Sup. (4) **26** (1993), no. 4, 489–516
[3] Y. Eliashberg and W. Thurston, *Confoliations*, Univ. Lect. Ser. **13**, AMS, Providence, RI, 1998.
[4] D. Gabai, *Foliations and the topology of 3-manifolds*, J. Diff. Geom. **18** (1983), no. 3, 445–503
[5] É. Ghys, *Laminations par surfaces de Riemann*, Panoramas et Syntheses
[6] N. Kopell, *Commuting diffeomorphisms*, in 1970 Global Analysis (Proc. Sympos. Pure Math., Vol. XIV, Berkeley, Calif., 1968), pp. 165–184
[7] A. Navas, *Groups of Circle Diffeomorphisms*, Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 2011

[8] S. Novikov, *Topology of Foliations*, Trud. Mosc. Mat. Obs. **14** (1965), 248–278
[9] S. Novikov, *Topology of Foliations*, English translation of [8]; available at https://homepage.mi-ras.ru/~snovikov/23.pdf
[10] J. Plante, *Foliations with measure preserving holonomy*, Ann. of Math. (2) **102** (1975), no. 2, 327–361
[11] R. Sacksteder, *Foliations and pseudogroups*, Amer. J. Math. **87** (1965), 79–102
[12] M. Scharlemann, *Sutured manifolds and generalized Thurston norms*, J. Diff. Geom. **29** (1989), 557–614
[13] F. Sergeraert, *Feuilletages et difféomorphisms infiniment tangents à l'identité*, Invent. Math. **39** (1977), 253–275
[14] W. Sierpinski, *Un theoreme sur les continus*, Tohoku Math. J. **13** (1918), 300–303
[15] 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