

Notes from U. Texas at Austin, Spring 1993

Mark Brittenham

Foliations and the Topology of 3-manifolds

Outline of class 1

Ultimately we will be focussing on codimension-1 foliations of 3-manifolds, and, more specifically, on what such a foliation can tell us about the topology of the 3-manifold that we have foliated. But we will begin with some generalities, aimed at familiarizing ourselves with various ways of thinking about foliations, and some of the basic concepts used in manipulating them.

Some source materials:

Survey articles:

Lawson, *Foliations*, Bull. AMS (May, 1974), 50 pp.

Gabai, *Foliations and 3-manifolds*, Proc. ICM Kyoto (1990), 10 pp.

Lecture notes:

Conlon, *Foliations of codimension-1*, year-long course at Wash. U., 1989-90, 300 pp.
(Available for photocopying from the instructor)

Books:

Tamura, *Topology of foliations*

Smalchuk & Net, Geometric Theory of Foliations
(instructor has a photocopy)

Original sources:

Novikov, *Topology of foliations*, Moscow Math. Soc. (Trudy) (1965), 50 pp.

Gabai/Thurston, *Genera of the arborescent links, A norm for the homology of 3-manifolds*, Mem. AMS **339** (Jan., 1986), 130 pp.

First definition (topologists' definition):

A **codimension-k foliation** \mathcal{F} of an n -manifold M is a way of decomposing M into a collection of (disjoint images of 1-to-1 immersions of) path-connected $(n-k)$ -dimensional manifolds, so that locally the manifolds look like the horizontal leaves of a product:

for all $x \in M$, there is an open neighborhood \mathcal{U} of x and a homeomorphism $h: \mathcal{U} \rightarrow (-1,1)^n$, so that the $(n-k)$ -manifolds get mapped under h to horizontal sheets

$$(-1,1) \times \dots \times (-1,1) \times \{x_{n-k+1}\} \times \dots \times \{x_n\}$$

The idea is that if you don't look too closely, you'd think you were looking at the fibers of an $(n-k)$ - dimensional vector bundle.

Each path-connected $(n-k)$ - manifold of the decomposition is called a **leaf** of the foliation.

Examples:

A vector bundle over a k -dimensional manifold - the fibers give a codimension- k foliation of the total space. Proof: local triviality of the bundle.

More generally, a fiber bundle over a k -manifold with fiber a manifold gives a codimension - k foliation (by fibers) of the total space.

Still more generally, a submersion $f: M^n \rightarrow N^k$ between smooth manifolds (i.e. $f_* : T_x M \rightarrow T_{f(x)} N$ is a surjection for all $x \in M$) gives rise to a codimension - k foliation of M by (path-components of) point inverses $f^{-1}(*)$. Proof: the Implicit Function Theorem says that by the appropriate choice of local coordinates the function f , at any point, looks like projection onto the last k coordinates: $f(x_1, \dots, x_n) = (x_{n-k+1}, \dots, x_n)$. Therefore, point-inverses look, locally, like horizontal sheets. A posteriori, each leaf is an $(n-k)$ - manifold.

So, for example, $f: \mathbf{R}^3 \rightarrow \mathbf{R}^1$ given by $f(x,y,z) = x^7 + xz^6 + y^9z + y^2 + 3z$ is a submersion (just check that $\nabla f \neq 0$ everywhere).

Exercise: Show that $f: \mathbf{R}^3 \rightarrow \mathbf{R}^1$ given by $f(x,y,z) = (x^2 + y^2 - 1)e^z$ is a submersion, so foliates \mathbf{R}^3 . What do the leaves look like? (Note: they will be symmetric about the z -axis; there is a qualitative difference between negative, positive, and zero-values).

Note: every leaf of the induced foliation is a closed subset of M (i.e., the 1-to-1 immersions are **embeddings**) - this is because each point-inverse is a closed set, and each path component is an open subset of the point-inverse (every point has a path-connected neighborhood), so the **complement** of a component is a union of path components, hence is open in the point-inverse.

Fact: for submersions from Euclidean spaces to Euclidean spaces, **none** of the leaves are compact (proof later - uses the notion of **holonomy**).

For a generic map $f:M \rightarrow N$, one can simply delete the set of points $\Sigma(f)$ where f_* is not surjective, leaving an open subset of M , by the Implicit Function Theorem), and so obtain a submersion $f:M \setminus \Sigma(f) \rightarrow N$, which yields a foliation away from the singular set. For example, $f:\mathbf{R}^3 \rightarrow \mathbf{R}^1$ given by $f(x,y,z) = xyz$ foliates $M=\mathbf{R}^3 \setminus (\text{the coordinate axes})$ by planes.


Completely unrelated question: what is the fundamental group of M ?

There are two other topological spaces (which can really be thought of as different topologies on the set M) which a foliation \mathcal{F} of M gives rise to. The first is the space of leaves.


Put an equivalence relation on M by $x \sim y$ if x and y are in the same leaf of \mathcal{F} (it should be easy to see that this is an equivalence). There is a natural surjection $p:M \rightarrow M/\sim (= (\text{notation}) M/\mathcal{F})$ (send a point to its equivalence class), and we can give M/\mathcal{F} the quotient topology induced by this map ($\mathcal{U} \subseteq M/\mathcal{F}$ is open iff $p^{-1}(\mathcal{U})$ is open in M). This is the space of leaves.

Example: Let \mathcal{F} be the foliation of \mathbf{R}^2 pictured below (there is actually a homeomorphism of \mathbf{R}^2 taking it to the foliation on the right). Then the space of leaves can be identified in pieces - the two collections of horizontal planes (in the right-hand picture) descend to closed half-rays, while the leaves in between descend to an open half-ray, which have both endpoints of the closed rays as limit points of the open end of the ray - see the picture. Each point of the space of leaves, it is easy to see, has a neighborhood homeomorphic to \mathbf{R}^1 , so it is a (non-Hausdorff) 1-manifold (the two endpoints cannot be separated).

Note: this example actually arises as a foliation induced by a submersion $f: \mathbf{R}^2 \rightarrow \mathbf{R}^1$; one can assign a function which is constant on leaves, which has a nowhere-zero gradient (just start at $-\infty$ and work up to 0 (without getting there) on the open ray, and start at 0 and go to ∞ on each of the closed rays - see the Figure).



Exercise: Build foliations of \mathbf{R}^2 with the spaces of leaves given below. Can you build a submersion inducing them? What kinds of limiting behaviors can we assign to the ends (tending to $-\infty$ or ∞)?



Some facts about foliations of \mathbf{R}^2 :

1. Every leaf of every foliation of \mathbf{R}^2 is a plane.
2. The space of leaves is always a (usually non-Hausdorff) 1-manifold. It always turns out to have trivial fundamental group!

We will eventually prove both of these statements, when we have developed the proper tools.

3. Every simply-connected (non-Hausdorff) 1-manifold is the space of leaves of some foliation of \mathbf{R}^2 . (Haefliger and Reeb, 1958)

4. Every foliation of \mathbf{R}^2 come from a submersion $f:\mathbf{R}^2 \rightarrow \mathbf{R}^1$. (Kaplan, 1940)

I doubt we will prove these two - they involve a fair amount of analysis. Besides, it's practically impossible to understand mathematics that was written in 1940.

The second space is the leaf topology on M - we'll cover that next time.

Foliations and the Topology of 3-manifolds

Outline of class 2

The second space commonly associated to a codim - k foliation \mathcal{F} of a manifold M is the leaf topology (denoted $M_{\mathcal{F}}$) on M ; basically, it is a way of putting a (finer) topology on M so that it becomes homeomorphic to a disjoint union of its leaves. A basis for the topology consists of open (in \mathbf{R}^{n-k}) subsets of all of the **plaques** of the foliation (i.e., the horizontal sheets in each of the distinguished coordinate charts of the definition).

Note: you should check that this is in fact a basis for a topology - a proof I think requires Invariance of Domain (in dimension $n-k$).

This topology is finer than the usual manifold topology on M , since any open subset of M can be written as a union of open subsets of distinguished coord. charts, each of which in turn is a union of open subsets of plaques. so the identity map $I: M_{\mathcal{F}} \rightarrow M$ is a continuous function.

We should also note that a single leaf can meet a given distinguished coordinate chart in only a countable number of plaques, since the 'centers' of each plaque would form a discrete set in the (manifold topology of the) leaf, and a path-connected manifold cannot contain an uncountable discrete set (exercise). From this it is easy to conclude that every foliation (except codimension-0 ones!) has an uncountable number of leaves, since any coordinate chart has an uncountable number of plaques. Thus it is easy to see that the idea of a path in a leaf (the continuous image of an interval) is the same in both the subspace topology (that the image of the leaf inherits from the usual topology on M) and the leaf topology on M ; a path could not move 'transversely' in a chart (by the Intermediate Value Theorem!). This same idea of counting transversely (or what really amounts to projecting charts onto the \mathbf{R}^k 's spanned by the last k coordinates of \mathbf{R}^n) makes it easy to see that the identity function above takes (sequentially) compact sets to (sequentially) compact sets (since a compact set in $M_{\mathcal{F}}$ can meet only finitely-many plaques in any coordinate chart). (Never mind - that's obvious! The identity function is continuous!)

This point of view of a foliation, as a way of cutting up a manifold into lower-

dimensional objects that fit together nicely, will be our main point of view for most of the things we will want to be doing with them. But there are other points of view that can also be useful at times, so we will spend some time exploring them, too.

Our second point of view is that a foliation is some kind of 'geometric structure' on the manifold.

A manifold is just a space M with, for each point $x \in M$, an open set \mathcal{U} and a homeomorphism $h: \mathcal{U} \rightarrow \mathbf{R}^n$ (which we call a coordinate chart, since it is a way of imposing a coordinate system on \mathcal{U} , by pulling back the one on \mathbf{R}^n). The \mathcal{U} 's cover M , and therefore give us an **atlas** of coordinate charts on M , $\{(\mathcal{U}_\alpha, h_\alpha)\}$.

Extra structures are imposed on M by imposing extra conditions on the **transition functions**

$$t_{\alpha\beta} = h_\beta \circ h_\alpha^{-1}: h_\alpha(\mathcal{U}_\alpha \cap \mathcal{U}_\beta) \rightarrow h_\beta(\mathcal{U}_\alpha \cap \mathcal{U}_\beta)$$

of the atlas. For example, if we require that the transition functions (which map open subsets of \mathbf{R}^n to open subsets of \mathbf{R}^n) are all C^r -smooth, then we have a C^r -manifold. If we require that they are all piecewise linear (PL), then we have a PL-manifold. If we require that they are all real-analytic (denoted C^ω), then we have a C^ω -manifold. If we only require homeomorphisms, we have a TOP-manifold. Many more examples can be easily found, e.g., if transition functions are Euclidean isometries, we have a Euclidean (or flat) manifold; if they are isometries of hyperbolic space (thinking of \mathbf{R}^n as hyperbolic n -space), we have a hyperbolic manifold.

Foliations can be thought of in much the same way, as a 'reduction of the structure group' of the manifold, since having a foliation imposes restrictions on the kind of transition functions we have. This is easy to see if we imagine two coordinate charts intersecting one another. On the intersection we have two maps to \mathbf{R}^n carrying the leaves of the foliation to horizontal sheets. Therefore the transition function between them carries a horizontal sheet to a leaf of the foliation to a horizontal sheet; i.e., the transition functions preserve the horizontal sheets $\mathbf{R}^{n-k} \times \{\text{pt.}\}$ of \mathbf{R}^n . Put slightly differently:

Second definition (differential topologists' definition):

A C^r codimension- k foliation of an n -manifold M is an atlas of coordinate charts for M whose transition functions

$$t_{\alpha\beta} = (t_{\alpha\beta}^1, \dots, t_{\alpha\beta}^n)$$

are C^r and have the property that $t_{\alpha\beta}^{n-k+1}, \dots, t_{\alpha\beta}^n$ (which a priori are functions of x_1, \dots, x_n) are functions only of x_{n-k+1}, \dots, x_n .

This is because what sheet you are in depends only on the last k coordinates, so if sheets get carried to sheets, what the last k coordinates of the transition function are depends only on the last k coordinates of the point we're at.

Seeing that a foliation according to the old definition give a foliation according to the new one is basically just the argument we gave above; going the other way, we have our distinguished coordinate charts, and the (immersed images of the) submanifolds can be just pieced together from the (inverse images of the) horizontal sheets by some sort of 'analytic continuation'.

This point of view is very useful in some circumstances, since it makes a foliation seem alot like other 'geometric structures' that have been the focus of a great deal of research (reductions of structure groups can be related to lifts of 'classifying maps' from one classifying space to another). In fact some people have used foliations as a more amenable 'testing ground' for techniques in this field.

One use we can put this point of view to is to help us understand how to pass foliations up and down a covering space projection. Suppose we are given a covering space projection $p: \tilde{M} \rightarrow M$, and, first, a codimension- k foliation \mathcal{F} of the base M . So around each point of M we have a distinguished coordinate chart for M , \mathcal{U} , and a homeomorphism $h: \mathcal{U} \rightarrow \mathbf{R}^n$. By the lifting property of covering spaces (since \mathcal{U} is simply-connected), \mathcal{U} is evenly-covered by p ; $p^{-1}(\mathcal{U}) =$ a disjoint union of sets \mathcal{U}_α each mapped homeomorphically down to \mathcal{U} by p . so if we take the collection of pairs $\{(\mathcal{U}_\alpha, \text{hop}|_{\mathcal{U}_\alpha})\}$ ranging over all preimages of (domains of) distinguished coordinate charts, we get an atlas on \tilde{M} , which, we claim, is a foliation. This is because we can easily check that the transition functions for these charts preserve horizontal sheets: if

$$(\mathcal{U}_\alpha, \text{hop}|_{\mathcal{U}_\alpha}) \text{ and } (\mathcal{V}_\beta, \text{kop}|_{\mathcal{V}_\beta})$$

are two such charts (with intersecting domains), then their transition function is

$$(\text{hop}|_{\mathcal{U}_\alpha}) \circ (\text{kop}|_{\mathcal{V}_\beta})^{-1} = \text{ho}(p|_{\mathcal{U}_\alpha} \circ (p|_{\mathcal{V}_\beta})^{-1}) \circ k^{-1} = \text{hok}^{-1}$$

(since going up from the base and then down again gives the identity), which preserves horizontal levels. It is not hard to see, in fact, that the foliation of \tilde{M} that this creates is the one whose leaves are (path-components of) the inverse images of the leaves of \mathcal{F} downstairs; think analytic continuation again.

This technique can build interesting new foliations for us (when the manifolds we know how to foliate have non-trivial fundamental group, which unfortunately is not the case for nearly every foliation we have built so far!). But doing this the other way around is even more useful.


Suppose we now have our covering space and a foliation $\tilde{\mathcal{F}}$ of \tilde{M} . Let us figure out what it takes for this foliation to descend to a foliation of M . It's easy to imagine what we would like to do - namely, to push the distinguished coordinate charts down onto M (i.e., to compose them with the inverse of p , restricted to a small enough set). But it is easy to see that, unlike the previous case, this will not always work - a set downstairs may have many lifts upstairs, and the foliations on them upstairs may not project 'compatibly'. In fact, what we need is that, for any distinguished coordinate chart downstairs, and any two lifts of it $(\mathcal{U}_\alpha, h), (\mathcal{U}_\beta, k)$ upstairs, the composition

$$(h \circ p|_{\mathcal{U}_\alpha}^{-1}) \circ (k \circ p|_{\mathcal{U}_\beta}^{-1})^{-1} = h \circ (p|_{\mathcal{U}_\alpha}^{-1}) \circ p|_{\mathcal{U}_\beta} \circ k^{-1}$$

need not preserve horizontal leaves, since going down then up need not be the identity. (The reasoning here is just slightly backwards, but I hope it conveys the correct idea). Basically what is needed is that $p|_{\mathcal{U}_\alpha}^{-1} \circ p|_{\mathcal{U}_\beta}$ send (pieces of) leaves of $\tilde{\mathcal{F}}$ to leaves of $\tilde{\mathcal{F}}$. This must in general be (laboriously) checked for every pair of lifts - that going down then back up sends leaves to leaves.

But we can do much better in some cases. Instead of going down from \mathcal{U}_β and then back up to \mathcal{U}_α , we can go directly from one to the other by a map which makes the diagram below commute. Such a map (if it extends to a homeomorphism of \tilde{M} to itself, which commutes with the projection) is called a covering translation (or deck transformation)

of \tilde{M} . Since this map then preserves the pieces of the leaves of the foliation $\tilde{\mathcal{F}}$, it then must in fact send leaves of $\tilde{\mathcal{F}}$ to leaves of $\tilde{\mathcal{F}}$, i.e., it leaves the foliation invariant (Exercise: convince yourself that this works the other way around, too).



But we need such maps for every pair of lifts, so we need the deck transformation group (the set of all deck transformations) to act transitively on lifts. A covering space with that property is called regular (or normal), and a necessary and sufficient condition that a covering be normal is for $p_*\pi_1(\tilde{M})$ to be a normal subgroup of $\pi_1(M)$.

The easiest way to arrange that a covering space is normal is to have $\pi_1(\tilde{M}) = \{1\}$ (the trivial subgroup is normal), i.e., \tilde{M} is the universal covering of M . Then one only needs to check that the foliation $\tilde{\mathcal{F}}$ on \tilde{M} is invariant under the covering translations of p , to insure that $\tilde{\mathcal{F}}$ descends to a foliation of M . Exercise: the leaves of the resulting foliation are the images of the leaves of $\tilde{\mathcal{F}}$!

As an example, the 2-torus $T^2 = S^1 \times S^1$ has universal cover \mathbf{R}^2 , and the group of covering translations are literally translations of \mathbf{R}^2 by integer quantities in x - and y -directions; $G = \{f : f(x,y) = (x+n, y+m) \text{ for some } n, m \in \mathbf{Z}\}$. If we foliate \mathbf{R}^2 by parallel lines of some fixed slope a/b , $\tilde{\mathcal{F}} = \{(x,y) : ax+by=c\} : c \in \mathbf{R}\}$, then it is easy to see that this foliation is invariant under any element of G , since if $ax+by=c$, then $a(x+m)+b(y+n)=c+an+bm$ (i.e., lines of slope a/b are carried to lines of slope a/b by translations!). So these foliations (there is one for each $a/b \in \mathbf{R} \cup \{\infty\}$) all descend to foliations of T^2 .

Similarly, if we foliate \mathbf{R}^3 by parallel planes $\{(x,y,z) : ax+by+cz=d\} : d \in \mathbf{R}\}$ (for fixed $a,b,c \in \mathbf{R}$), then these foliations are also invariant under translations of \mathbf{R}^3 by integer (and in fact any) amounts, so descend to foliations of the three-torus $T^3 = S^1 \times S^1 \times S^1$.

Next time we will explore a bit what these foliations look like, and then move on to describe a third way (the differential geometers' way) of thinking about a foliation.

Foliations and the Topology of 3-manifolds

Outline of class 3

Last time we saw that foliations of \mathbf{R}^2 and \mathbf{R}^3 by parallel lines and planes, respectively, descend to foliations of 2- and 3-tori. (In fact, if you foliate \mathbf{R}^n by parallel hyperplanes (of codimension- k), it descends to a codimension- k foliation of the n -torus T^n . We'll begin by studying the limiting behaviors of the leaves of these foliations.

For T^2 , there is a qualitative difference in the foliations coming from $\tilde{\mathcal{F}} = \{(x,y) : ax+by=c\} : c \in \mathbf{R}\}$, depending on whether a/b is rational or not. If $a/b = p/q$ in lowest terms (where $p,q \in \mathbf{Z}$), then it is easy to see that the translation $f(x,y)=(x+q,y-p)$ carries a leaf of the foliation onto itself (since $aq-bp=0$), so each leaf of $\tilde{\mathcal{F}}$ descends to a circle downstairs. So \mathcal{F} is a foliation all of whose leaves are circles. We'll see later that this will imply that the space of leaves of \mathcal{F} is a circle, and so this is actually a foliation by leaves of a circle bundle.

If $a/b \notin \mathbf{Q}$, then it is also easy to see that no covering translation carries a leaf onto itself (a point and its image would demonstrate that the leaf had rational slope), so every leaf of $\tilde{\mathcal{F}}$ maps injectively down to T^2 , so every leaf of \mathcal{F} is a line.

But even more, we can see without too much difficulty that every leaf of \mathcal{F} is dense in T^2 . One way to see this is to note that if we take the projection of a given leaf $\{(x,y) : ax+by=c\}$ upstairs and the circle downstairs that is the projection of the horizontal line $\{(x,y) : y=0\}$, then the former meets the latter in the set $A = \{an+c \pmod{1} : n \in \mathbf{Z}\} \subseteq S^1$. This set is dense in the circle; the points are distinct for distinct n (otherwise a must be rational!), so this is an infinite subset of the circle. Since the circle is compact, the set has a limit point, so for any $\epsilon > 0$, there are two points of the set within ϵ of one another. This implies (by subtracting) that there are numbers n_0, m_0 such that $|an_0 + m_0| = r_0 < \epsilon$. This in turn implies that the points $\{k(an_0)+c \pmod{1} : 1 \leq k \leq N\} \subseteq S^1$ (for some $N > 1/r_0$) consists of points travelling all the way around the circle at distance less than ϵ from its predecessor. So every point of the circle is within ϵ of a point of A , for every ϵ , i.e., the intersection of every leaf is dense in this circle. To show that the leaf is in fact dense in T^2 , just pick any point and drag it along in its leaf to this circle; find a point in the (dense)

leaf as close to that point of the circle as you like, and then drag both back to the first's original position to find a point in the (dense) leaf as close to your point as you want.

There are similar things that can be said about the foliations of T^3 that we have built. If the coefficients of $\{(x,y,z) : ax+by+cz=d\} : d \in \mathbf{R}\}$ satisfy $a/b, b/c, c/a \in \mathbf{Q}$, then each leaf of $\tilde{\mathcal{F}}$ descends to a torus in T^3 , giving a foliation of T^3 by tori (which, again, we will see is a foliation coming from a fiber bundle). If one of them is irrational, on the other hand, then every leaf of the foliation downstairs is dense (a similar argument to the one given above, using a torus downstairs instead of a circle, will work). If, further, the coefficients a, b, c are linearly independent over the integers (i.e., $an+bm+cr=0$ with $n, m, r \in \mathbf{Z}$ implies $n=m=r=0$), then no covering translation identifies points of any single leaf (check!), so every leaf projects injectively to T^3 , giving a foliation of T^3 all of whose leaves are planes. As an example, take $(a, b, c) = (1, \sqrt{2}, \pi)$ (or your favorite triple of rational number, irrational algebraic number, and transcendental number).

This last fact has a rather amazing converse:

Theorem(Rosenberg(C^2 case), Gabai(C^0 case): If a closed 3-manifold M admits a foliation by planes, then $M = T^3$.

We might take a stab at the proof of this theorem at some point, although it involves some heavy-duty group theory!

Distributions.

Suppose we have a foliation \mathcal{F} of M^n with smooth leaves, i.e., the 1-to-1 immersions of the leaves of \mathcal{F} are C^1 or better; another way to put it is that the first $(n-k)$ coordinates of the transition functions for the atlas of charts are smooth. Then we can associate to each point x of M the $(n-k)$ -plane in $T_x M$ tangent to the leaf passing through that point (which we will denote $T_x \mathcal{F}$; you can think of it as the image under f_* , where f is the 1-to-1 immersion of the leaf, of the tangent space of the leaf). If these $(n-k)$ -planes vary continuously with x , then we get an $(n-k)$ -dimensional sub-bundle of TM (the tangent space of M), which we will call the tangent space to the foliation, and denote by $T\mathcal{F}$.

An $(n-k)$ -dimensional subbundle of TM is also known as a codimension- k (or $(n-k)$ -dimensional) **distribution** (usually denoted Δ^{n-k}) on M . Therefore, a necessary condition

that a manifold M admit a codimension- k foliation is that it admit a codimension- k distribution.

This already allows us to show that some (closed) manifolds cannot admit foliations of various codimensions. For example, a 1-dimensional distribution is a (continuous) choice of 1-dimensional subspace of $T_x M$. This is almost a choice of a nowhere-zero vector field V on M ; certainly, a non-zero vector field determines a 1-dimensional distribution (just take the span of $V(x)$ in $T_x M$). However, by lifting a 1-dimensional distribution to a double cover \tilde{M} of M , if necessary, we can ‘orient’ our distribution, so that it is the span of a nowhere-zero vector field of M or \tilde{M} (details below). So if a manifold M admits a 1-dimensional distribution, then either M or a double cover \tilde{M} admits a nowhere-zero vector field; but it is a well-known fact that the Euler characteristic $\chi(M)$ of a manifold can be calculated by summing up the indices of zeros of a vector field on M . If V has no zeros, therefore, $\chi(M)=0$ (empty sum). So if M admits a 1-dimensional distribution, then $\chi(M)=0$ or $\chi(\tilde{M})=0$. But \tilde{M} is a double cover of M , so $\chi(\tilde{M})=2\chi(M)=0$, so $\chi(M)=0$. Since a (closed) manifold with $\chi(M)=0$ does in fact admit a nowhere-zero vector field (a mildly difficult exercise), we therefore have:

Prop.: A closed n -manifold admits a 1-dimensional distribution if and only if $\chi(M)=0$.

So, for example, no even-dimensional sphere S^{2n} admits a 1-dimensional foliation, since $\chi(S^{2n})=2$. Nor does any closed surface, except the torus and the Klein bottle, since all others have non-zero Euler characteristic.

The construction of the two-fold cover required in the argument above follows a standard line, and works equally well anytime you are trying to ‘orient’ a distribution. The idea is to ‘find’ the two-fold cover by instead finding (the image of) its fundamental group. I.e., we instead find an index 2 subgroup of $\pi_1(M)$, and appeal to covering space theory to tell us that it corresponds to a 2-fold cover of M . We do this by picking a basepoint x and an orientation for the fiber Δ_x^1 at that point. Then we try to drag it around loops to see if we come back with the same or opposite orientation. The idea is that we can give the nearby fibers an orientation consistent with the orientation on any one fiber (this should become clear soon), so we can pass the orientation along the loop, a little bit at a time,

until we come back to the beginning. We will let $G \subseteq \pi_1(M) = \pi_1(M, x)$ denote the set of homotopy classes of loops which have a representative which, when you drag the orientation on Δ_x^1 around it, returns with the same orientation; the ‘orientation-preserving’ loops of M . It is easy to see that this property in fact depends only on the homotopy class of the curve, since any homotopy can be imagined as a sequence of really ‘small’ homotopies (for which the image of the support is contained in a neighborhood we can coherently orient), and since it is easy to see that the property of being orientation-preserving is preserved under small homotopies, it is preserved under all homotopies. It is also easy to see that G is a subgroup of $\pi_1(M)$, since the concatenation of two orientation-preserving loops is orientation-preserving (so G is closed under multiplication), and dragging an orientation around an orientation-preserving loop in the opposite direction will preserve the orientation (so G is closed under inversion). It is also easy to see that G has at most 2 cosets in $\pi_1(M)$; if there are no orientation-reversing loops, then $G = \pi_1(M)$, while if there is one γ , then any other γ' is equal to $\gamma \cdot \alpha$, where $\alpha \in G$ (i.e., $\alpha = \gamma' \cdot \gamma^{-1}$ is orientation-preserving; the first loop reverses orientation, and the second one reverses it back!). Consequently, G and γG are the only two cosets, so G has index 2.

Then we can lift the distribution Δ^1 to a distribution $\tilde{\Delta}^1$ on the 2-fold (or 1-fold!) cover \tilde{M} corresponding to G (just take the preimage of Δ_x^1 under the map $p_* : T_{\tilde{y}}\tilde{M} \rightarrow T_y M$ for each $\tilde{y} \in \tilde{M}$). This distribution can then be oriented by picking a base point \tilde{x} and an orientation for $\tilde{\Delta}_x^1$ and then assigning an orientation for any other $\tilde{\Delta}_y^1$ by dragging an orientation along any path from \tilde{x} to \tilde{y} . This assignment is independent of the path, since any two paths to the same point together form a loop γ . If paths assign different (i.e., opposite) orientations at \tilde{y} , then going around the loop drags the orientation at the basepoint to its opposite. But then its projection downstairs is also an orientation-reversing loop (dragging the orientation around, you wouldn’t be able to tell if you were upstairs or down, because it is a local construction). But this is absurd, since $p_*(\gamma) \in G$, which consists of orientation-preserving loops.

This therefore gives us the required construction of an orientable 2-fold cover. Notice that this immediately implies that any 1-dimensional distribution on a simply-connected manifold M must be orientable; $\pi_1(M) = \{1\}$ has no index-2 subgroup, so $\tilde{M} = M$!

Foliations and the Topology of 3-manifolds

Outline of class 4

Given a Riemannian metric on a manifold M (a choice of positive definite inner product on $T_x M$ for each $x \in M$) and a codimension- k distribution Δ on M , we can, by taking the orthogonal complement of Δ_x (which is a k -dimensional subspace of $T_x M$) find a k -dimensional distribution on M . Therefore, our result from last time, that a closed manifold M has a 1-dimensional distribution iff $\chi(M)=0$, implies also that a closed manifold M has a codimension-1 distribution iff $\chi(M)=0$. So manifolds with non-zero Euler characteristic (like the ones described last time) cannot admit a codimension-1 foliation.

But what can we say if a manifold does have a codimension- k distribution? There are several questions we might ask about foliations, that relate to distributions:

1. If M has a codimension- k distribution Δ , is $\Delta = T\mathcal{F}$ for some codimension- k foliation \mathcal{F} of M ?
2. (somewhat weaker) If M admits a codimension- k distribution, does it admit a codimension- k foliation?
3. Can the same distribution be the tangent space to two different foliations?

Two of these questions, at least, can be answered by the point of view offered by our third definition. The point is that there is a succinct criterion that must be satisfied for a distribution to be tangent to a foliation.

Definition: A codimension- k distribution Δ is **integrable** if for any two vector fields X, Y with $X(p), Y(p) \in \Delta_p$ for all $p \in M$, then the **Lie bracket** of X and Y , denoted $[X, Y]$, satisfies $[X, Y](p) \in \Delta_p$ for all $p \in M$.

Then the main theorem is:

Frobenius' Theorem: A distribution is the tangent space to a foliation if and only if it is integrable.

For an interpretation of the Lie bracket of two vector fields (another will be given later), we need a different interpretation of what a tangent vector is. The idea is that, given a (smooth) function f on M , we can talk about the directional derivative $D_v(f)$ of f in the

direction of a tangent vector, and this assignment is a derivation: $D_v(fg) = fD_v(g) + gD_v(f)$. This point of view can be turned around, however: we could in fact define a tangent vector (field) as a derivation which assigns to each smooth function another function (its directional derivative). Then the Lie bracket of two vector fields is the derivation which assigns to each function f the function $[X, Y](f) = X(Y(f)) - Y(X(f))$ (a quick calculation will convince you that this is in fact a derivation). The basic idea is that the Lie bracket measures the extent to which mixed partial derivatives fail to commute.

The idea of the proof of Frobenius' Theorem is that integrability is exactly the criterion it takes to find, around each point of M , a coordinate chart (h, \mathcal{U}) which carries $\Delta|_{\mathcal{U}}$, under h_* , to the horizontal distribution, i.e., the set of subspaces whose vectors have last k coordinates zero.

It is easy to see that the horizontal distribution is the tangent space to a foliation on \mathbf{R}^n , namely the horizontal foliation (the set of $(n-k)$ -planes obtained by setting the last k coordinates to constants). To help us with the last question on our list (as well as our theorem), we should notice that the horizontal foliation is the only codimension- k foliation of \mathbf{R}^n tangent to the horizontal distribution. Because if there were another foliation, then there would be a leaf containing points at different horizontal levels. So if we take a path in that leaf between the points (which therefore has tangent vector (tangent to the leaf so) in the distribution) and then project it onto the coordinate where the two points differ, we would get a smooth function from an interval to \mathbf{R} . But then the Mean Value Theorem implies that some where in between the function has non-zero derivative. But this means that at that point of the path its tangent vector has a vertical component, a contradiction.

Therefore, if we have two overlapping such coordinate charts (h, \mathcal{U}) , (k, \mathcal{V}) . then by pulling back the horizontal foliation on one, we get a foliation on \mathcal{U} with tangent space $\Delta|_{\mathcal{U}}$, and then pushing it forward under k we get a foliation on \mathbf{R}^n which is tangent to the horizontal distribution, and hence (by the previous paragraph) is the horizontal foliation. So the transition function takes horizontal sheets to horizontal sheets, so is a foliation.

We can now quickly answer the third question - the answer is 'No' - since if there were two foliations with the same tangent space, then pushing each forward under a dis-

tistinguished coordinate chart, we get two foliations on \mathbf{R}^n both tangent to the horizontal distribution, a contradiction.

Also, this point of view allows us to quickly see that every 1-dimensional distribution is the tangent space to a foliation. This is because any two vector fields X, Y with $X(p), Y(p) \in \Delta_p^1$ for all p are, locally, multiples of some non-zero vector field Z (so $X = aZ$, $Y = bZ$ for some functions $a, b: \mathcal{U} \rightarrow \mathbf{R}$), and then a quick calculation shows

$$\begin{aligned} [X, Y](f) &= aZ(bZ(f)) - bZ(aZ(f)) = a(Z(b)Z(f) + bZ(Z(f))) - b(Z(a)Z(f) + aZ(Z(f))) \\ &= (aZ(b) - bZ(a))Z(f) \end{aligned}$$

so $[X, Y](p) = (aZ(b) - bZ(a))Z(p) \in \Delta_p^1$ for all $p \in M$, i.e. Δ^1 is integrable. This fact is actually the existence and uniqueness of solutions to ordinary differential equations.

Integrability in higher dimensions is a non-trivial criterion, however (so the answer to our first question is also 'No'). Examples are not hard to come by; for example, on \mathbf{R}^3 we can take the 2-dimensional distribution

$$\Delta_{(a,b,c)}^2 = \text{span of } \{(0,1,0), (1,0,b)\}$$

We will show that this is not integrable by showing that if it were the tangent space to a foliation, then the xy -plane must be a leaf. But this is impossible, since its tangent space at $(1,1,0)$, would then be the xy -plane, which is not the same as $\Delta_{(1,1,0)}^2 = \text{span of } \{(0,1,0), (1,0,1)\}$, since the latter contains vectors with non-zero z -coordinate.

The idea behind this is that if $\Delta^2 = T\mathcal{F}$ and γ is a path in M with $\gamma'(t) \in \Delta_{\gamma(t)}^2$ for all t , then γ is contained in a single leaf of \mathcal{F} . This follows from work above: if we push a piece of the path forward under a distinguished coordinate chart, we get a path in \mathbf{R}^3 whose tangent vectors are in the pushed-forward (i.e. horizontal) distribution, so the argument above shows that the path is horizontal, i.e., is entirely contained in a horizontal leaf. So the piece of the path back in M is entirely contained in a leaf of \mathcal{F} . For this it is easy to see that the set of all points t such that $\gamma(t)$ is in a given leaf is open; so if γ is not contained in a single leaf, then the unit interval can be written as the union of two or more disjoint open sets, so it would not be connected, a contradiction. (Note that this argument is completely general, and applies to the tangent space to any foliation.)

If we use this on the distribution given above, assuming that it is tangent to a foliation, then since (with $\gamma(t)=(t,0,0)$) $\gamma'(t)=(1,0,0)\in\Delta_{(t,0,0)}^2$ for all t , so the image of γ (i.e., the x -axis), would be entirely contained in a leaf. But then for each $a\in\mathbf{R}$ (with $\gamma(t)=(a,t,0)$) $\gamma'(t)=(0,1,0)\in\Delta_{(a,t,0)}^2$ for all t , so each would also be contained in a single leaf. But since these all intersect the x -axis, which is contained in a single leaf, each would be contained in that same single leaf, so their union, the xy -plane, would be contained in a single leaf of the assumed foliation. But the only way this could be true is if the xy -plane is a leaf.

Given a codimension-1 foliation of a manifold M , we've seen that using a Riemannian metric we can construct an (orthogonal) 1-dimensional distribution on M , which we now know is integrable. So every codimension-1 foliation \mathcal{F} admits a **transverse** 1-dimensional foliation (which means that at each point of M , their tangent spaces together span the tangent space of M). Recalling our construction from last time, we can, by possibly passing to a 2-fold cover of M , orient the transverse foliation (i.e., orient its tangent space). We can therefore lift our codimension-1 foliation to one whose transverse foliation is orientable. Such a codimension-1 foliation is called (not surprisingly) **transversely orientable**. Since the 1-dimensional distribution is the span of a vector field, this also means that the codimension-1 foliation has a transverse vector field. This will be a very useful construction for us later, since we will be able to use this to 'push' objects and constructions in a leaf orthogonally off of the leaf, and lift them to nearby leaves.

Foliations and the Topology of 3-manifolds

Outline of class 5


We have seen that a (closed) manifold can have a codimension-1 foliation only if its Euler characteristic is zero. We have also seen that we can transversely orient a codimension-1 foliation, after perhaps passing to a 2-fold cover of our manifold. This immediately implies that a codimension-1 foliation of a simply-connected manifold is transversely oriented. These two facts together allow us to prove two of the facts about foliations \mathcal{F} of \mathbf{R}^2 stated in the first lecture:

1. Every leaf of the foliation is a line.

The only other alternative is that some leaf is a circle, γ . But then by the Jordan Curve Theorem, γ bounds a disk Δ^2 in \mathbf{R}^2 . This disk inherits a codimension-1 foliation from \mathbf{R}^2 , with $\gamma = \partial\Delta^2$ as a leaf. But if we take two copies Δ_0^2, Δ_1^2 of this foliated disk and glue them together along their boundaries, we have a 2-sphere S^2 with a codimension-1 foliation, a contradiction, since $\chi(S^2)=2$.

2. The space of leaves of \mathcal{F} is a (non-Hausdorff) manifold.

We will show this by proving the **claim**: no leaf of \mathcal{F} meets a distinguished chart for the foliation twice. This implies the result, since then an open arc tranverse to a distinguished chart maps injectively into the space of leaves, and maps to an open set (giving a locally-Euclidean neighborhood for every point in it). It is an open set since its inverse image in \mathbf{R}^2 , which is the union of all leaves which intersect the arc, is open in \mathbf{R}^2 (for any point in the inverse image, draw a path in its leaf to the arc, then it's fairly easy to see the a neighborhood of the endpoint in the arc can be dragged back to get a neighborhood of our point entirely in the inverse image (see the figure below).



To prove the claim, suppose there were a leaf hitting a chart twice. This gives a path in a leaf joining two points in different levels together. \mathbf{R}^2 is simply connected, so our foliation is transversely oriented, so there is a vector field transverse to the foliation. WOLOG we can assume that we orient the path in the leaf so that the starting point is 'above' the ending point in the chart (in the sense of the transverse vector field). We can then use the vector field to flow the path slightly off of itself (see figure below), giving

instead a path transverse to the foliation. Since the ends are in the same chart, we can join them together by a positively-oriented transverse arc to complete our pushed-off path to a (simple) loop γ everywhere transverse to the foliation \mathcal{F} . This loop again bounds a disk Δ^2 which inherits a foliation from \mathbf{R}^2 , except this time this foliation is everywhere transverse to the boundary $\partial\Delta^2=\gamma$. After ‘bending’ this foliation to be orthogonal along the boundary, we can again glue two of them together to form a foliation of the 2-sphere, a contradiction. So a leaf can’t meet a chart twice.

We have yet to answer the second question which we raised relating distributions to foliations:

If a manifold admits a codimension- k distribution, does it admit a codimension- k foliation?

If we think of a distribution as a section of the Grassman bundle, then we can talk of deforming a distribution - just homotope the section through sections. This really just means deforming the $(n-k)$ -planes at each point in a continuous fashion. Then we can ask an even stronger question:

Can we deform every codimension- k distribution to the tangent space to some foliation?

This stronger question was answered by Thurston in the mid-1970’s, with the surprising answer of ‘Yes’!

Theorem (Thurston): Any codimension-1 distribution on a closed manifold can be deformed to the tangent space of a C^∞ foliation of M . In particular, a closed manifold M admits a codimension-1 foliation if and only if $\chi(M)=0$.

Theorem: (Thurston) Any distribution of codimension greater than one can be deformed to the tangent space of a C^0 foliation (whose leaves are C^∞ - immersed).

The last theorem can be contrasted with the fact, well-known at the time Thurston proved these results, that there are non-trivial homotopy-theoretic obstructions (involving the cohomology of the normal bundle to the distribution) for a C^2 distribution to be C^2 -homotopic to the tangent space of a foliation. All of these last results are well outside of the scope of this course - we are about to focus our attention (for the remainder of the course) on codimension-1 foliations of 3-manifolds, where at least the weaker form of the first theorem has a much more accessible proof. That construction will be our next topic of study.

Constructing foliations on 3-manifolds

There is a classical (1960’s - for low-dimensional topology that is classical) construction of (transversely orientable) codimension-1 foliations of closed, orientable 3-manifolds, that we will work through now. We do so for two reasons - the construction introduces several ideas that will be useful later on, and, more importantly, the constructions will hopefully

make it clear what it is we want to avoid, if we want to make foliations topologically meaningful (which is our ultimate goal).

The basic idea of the construction is to use the (well-known, but we will sketch a proof here) fact that every closed, orientable 3-manifold M can be obtained by taking some link in the 3-sphere S^3 , removing solid torus neighborhoods of the loops, and gluing them back to the resulting link complement 'differently'. What we will do is show that we can build foliations of S^3 that can be carried across this construction to foliate M .

We will start by constructing the necessary foliations of S^3 . The main building block for this is the **Reeb component** or Reeb-foliated solid torus.

If we look at the foliation of \mathbf{R}^3 given by the submersion in the exercise of the first lecture, we see that the leaves are all axially symmetric about the z -axis (since what leaf a point is in depends only on $x^2 + y^2$, not on (x, y)). What the leaves look like breaks into 3 cases:

$(x^2 + y^2 - 1)e^z = 0$ implies $x^2 + y^2 = 1$, so the leaf is a vertical cylinder,


$(x^2 + y^2 - 1)e^z = -c^2 < 0$ implies $x^2 + y^2 = 1 - c^2 e^{-z} < 1$, so this leaf lives inside the cylinder; as $z \rightarrow \infty$, $x^2 + y^2 \rightarrow 1$, and as $z \rightarrow -\infty$ the right hand side becomes negative, so the leaf closes up by hitting the z -axis and does not continue on down,

$(x^2 + y^2 - 1)e^z = c^2 > 0$ implies $x^2 + y^2 = 1 + c^2 e^{-z} > 1$, so these leaves live outside the cylinder; as $z \rightarrow \infty$, $x^2 + y^2 \rightarrow 1$, and as $z \rightarrow -\infty$, $x^2 + y^2 \rightarrow \infty$, so these leaves flair out in the negative direction.

In other words, the leaves inside the cylinder are hyperboloids (so are planes), and the ones outside are annuli asymptotic to the cylinder in the positive direction, and flair out in the negative direction (see the figure below).

It is an easy matter to check that this foliation is invariant under vertical translations by integer (and in fact any) amounts. Since the vertical solid cylinder $D_2 \times \mathbf{R}^1$ is similarly invariant, its foliation descends to a foliation on the solid cylinder modulo these covering translations. This ambient space is the solid torus $D_2 \times S^1$. The cylinder leaf descends to the boundary torus, which is a leaf of the foliation. No \mathbf{R}^2 -leaf upstairs is carried to itself by a non-trivial translation (just look at the (one) point where each of these leaves hits the z -axis), they all descend to \mathbf{R}^2 -leaves downstairs. These leaves spiral out towards the torus leaf (this corresponds to travelling up in the z -direction), and so the foliation looks like the one pictured below. The core of the solid torus (the image of the z -axis) is a simple loop in $D_2 \times S^1$, which intersects each leaf in the interior of the solid torus transversely. It is called the **core** of the Reeb component.


It is also well-known that the 3-sphere S^3 can be obtained by gluing two solid tori together along their boundaries. The easiest way to see this is to think of the 3-sphere as two 3-balls glued together. If we drill out the neighborhood of an arc running straight from top to bottom, we turn the ball into a solid torus (see below); if we add this thickened arc to the other 3-ball we turn it into a solid torus as well, and these two solid tori meet along their boundaries. If we do this to the 'standard' 3-ball below (thinking of S^3 as $\mathbf{R}^3 \cup \{\infty\}$), we get the standard picture of these solid tori (the outside one contains the point at infinity). We can in fact do the same thing, drilling out g parallel arcs running vertically through one of the balls, turning it into a genus- g **handlebody**; adding the thickened arc to the other ball turns it into a handlebody as well, so we can write S^3 as the union of two genus- g handlebodies glued along their boundaries (this is known as a **Heegard decomposition**).



If we foliate each of the two solid tori as a Reeb component, then these foliations glue together nicely (the boundary torus is a leaf of both) to give a foliation of S^3 , known as the **Reeb foliation** of S^3 (if you're first, you get everything named after you). He supposedly built it in response to a recommendation by his advisor Ehresmann that a good thesis problem would be to show that S^3 admits no codimension-1 foliations! (It's likely, though, that Ehresmann was thinking of real analytic foliations (most of the early theory was focussed solely on them), which, it is true, the 3-sphere can't support - we'll see why later.) This foliation is our starting point.


We'll need some somewhat more complicated foliations on S^3 for our construction - they will all be built out of this one, though, by a process known as **turbulization**. To describe it we'll need to go back to the construction of the Reeb component. If, in the foliation of \mathbf{R}^3 above, we take a vertical solid cylinder, centered on the z -axis, which is slightly larger than the cylinder $x^2 + y^2 \leq 1$ that we took for the Reeb component, the foliation on it is also invariant under vertical translation, so it descends to a foliation on the solid torus as well. In this case however, the foliation upstairs is transverse to the boundary of the solid cylinder, so the foliation downstairs is transverse to the boundary

torus and meets it in (meridional) circles. Basically it looks like a Reeb-foliated solid torus which is strictly in the interior of our solid torus, together with half-infinite annuli which start at the boundary and spiral in towards the torus leaf (see below).



This foliation is central to the idea of turbulization. The idea is that if we have a foliation \mathcal{F} of M and a simple loop γ which is everywhere transverse to the leaves of \mathcal{F} , then by taking a small enough solid torus neighborhood of γ , we can insure that the foliation \mathcal{F} meets this solid torus in a collection of meridional disks (see below). In particular, it meets the torus boundary of this solid torus in (meridional) loops. If we erase this meridional foliation on the solid torus and replace it with the one we've constructed in the previous paragraph, the (because the foliation on the 'outside' of the solid torus meets the torus in meridional circles, which match up with the leaves inside the solid torus), we get a new foliation on M , and γ is now the core of a Reeb component of this new foliation (see below). The idea is that we have basically stuck the torus (the boundary of the Reeb component) into M tranverse to the foliation, around the loop γ , and have then 'spun' the leaves of the foliation as it approaches this torus. This new foliation is called the turbulization of \mathcal{F} about γ .

Next time we'll see how to use this idea of turbulization to foliate S^3 in a wide variety of ways, giving us foliations which we can 'carry across' the Dehn surgery construction, to foliate all closed orientable 3-manifolds.



Foliations and the Topology of 3-manifolds

Outline of class 6

We need two more ideas to complete our proof that all closed orientable 3-manifolds admit codimension-1 foliations. One comes from knot theory, and will be stated without proof.

Fact1: (Alexander) Every link in S^3 can be isotoped to a braid, i.e., it can be deformed into the standard solid torus in S^3 so that it runs everywhere transverse to the foliation of $S^1 \times D_2$ by meridional disks.

A proof may be found in Joan Birman's book, *Braids, links, and mapping class groups*. Alexander's original proof appeared in the 1920's in the Proceedings of the NAS.

If we instead replace the foliation of the solid torus by meridional disks by the Reeb foliation of the standard solid torus (where we either imagine that the braid lies very near the core of the solid torus, or the \mathbf{R}^2 leaves of the Reeb foliation don't start spinning until they are very near the boundary torus), then we can imagine instead that Alexander's result says that any link in S^3 can be isotoped to be transverse to the standard Reeb foliation of S^3 . Consequently, we can turbulize the Reeb foliation along such a link. So we get:

Proposition: For every link L in S^3 , there exists a foliation \mathcal{F} of S^3 such that each component of L is the core of a Reeb component of \mathcal{F} .

Notice that in fact $L \cup \{(z\text{-axis}) \cup \infty\}$ is a complete set of cores of the Reeb components of the turbulized foliation.

If we were to drill out a solid torus neighborhood of each component of L (getting the exterior of L , $E(L) \subseteq S^3$), and throw away the \mathbf{R}^2 -leaves of the Reeb components, we would in fact get a foliation of $E(L)$, such that the boundary tori are all leaves. This fact, together with the following one, will allow us to finish our construction.

Fact 2: For every closed, orientable 3-manifold M , there exists a link L in S^3 and a link L' in M such that $E(L)$ is homeomorphic to $E(L') = M \setminus \text{int}(N(L'))$.

This fact is usually stated differently, by saying that any (closed, orientable) 3-manifold can be obtained from S^3 by drilling out solid torus neighborhoods of some link, and then sewing the solid tori back in differently (this process is known as 'doing a Dehn surgery on the link'). It is easy to see that these two statements are equivalent.


We will outline a proof of this fact; but first let us use it to finish our construction. Given M , we find the promised links L in S^3 and L' in M . By the above argument, we can foliate $E(L)$, so since $E(L) \cong E(L')$, we can foliate $E(L')$ so that the boundary tori are leaves. But we can foliate $N(L)$ as a disjoint collection of Reeb solid tori, so by gluing the

pieces together, since the boundary tori are leaves of both foliations, is it easy to see that the two foliations together succeed in giving us a foliation of M .

The proof of Fact 2 proceeds in several steps, basically tracing its way through every standard way of picturing a 3-manifold. It starts with a classical result of Moise (first proved in the mid-1950's) that every 3-manifold possesses a triangulation, i.e., it can be written as a union of 3-simplices with pairs of faces glued together. Now if we take (what is known as a regular) neighborhood of the 1-skeleton of the triangulation, we get a handlebody (of some genus) H_g - it is a 3-ball (the regular neighborhood of a maximal tree in the 1-skeleton) with a bunch of 1-handles (the (relative) regular neighborhoods of the remaining edge) attached. Furthermore, the rest of M , $M \setminus \text{int}(H_g)$, is also a handlebody (of the same genus (since the genus of a handlebody is equal to the genus of its boundary) - it is a regular neighborhood of the 1-skeleton of the dual cell decomposition of M).

So every closed orientable 3-manifold can be thought of as two handlebodies of some genus (depending on M) glued together along their boundaries (by, technically, an orientation-reversing homeomorphism). But we have already seen that for any g , S^3 can be obtained by gluing two genus- g handlebodies together along their boundaries. So M can be obtained by splitting S^3 open along the genus- g decomposing surface and regluing the two handlebodies together by a different homeomorphism.

To get the two homeomorphisms in the same picture, let us say we get S^3 by gluing by the homeomorphism h , and M by gluing by the homeomorphism k . Let us instead imagine our manifolds as coming in three pieces, a handlebody on top and bottom, and a product $F_g \times I$ in between (see figure below). If we glue the bottom handlebody to the bottom of the product by h , then if we glue the top of the product to the top handlebody by the identity, we get S^3 ; but if we glue by the homeomorphism $k \circ h^{-1}$, we get M (since it undoes the first gluing by h). We will show the Dehn surgery picture by proving a fact about the homeomorphism $f = k \circ h^{-1}$ (which is an orientation preserving homeomorphism from F_g to itself).



Sub-fact: Every orientation-preserving homeomorphism of F_g is isotopic to a composition of a finite number of Dehn twists (about non-separating curves).

A **Dehn twist** homeomorphism D about an (oriented) simple closed curve γ is a homeomorphism which is the identity outside of a small annular neighborhood of γ , while in the annulus $S^1 \times [0,1]$ it is the homeomorphism $D(\theta, t) = (\theta \pm 2\pi t, t)$. $+$ gives a positive Dehn twist, and $-$ gives a negative Dehn twist.

The sub-fact is proved by induction on the genus g of F_g ; the base case, $F_0 = S^2$, due to Smale, says that every orientation-preserving homeomorphism of S^2 is isotopic to the identity. The induction step relies on



Claim: For any pair of non-separating simple curves γ, γ' there is a sequence D_1, \dots, D_n of Dehn twists along non-separating curves such that $D_n \circ \dots \circ D_1(\gamma')$ is isotopic to γ .

This fact allows us to carry out the inductive step, since if we pick any non-separating curve γ and set $\gamma' = f(\gamma)$, then there are Dehn twists so that $f_1(\gamma) = D_n \circ \dots \circ D_1 \circ f(\gamma)$ is isotopic to γ (so if we compose with an isotopy they are equal), and then if we split domain and range F_g along γ and glue in 2-disks we get a surface of genus $g-1$, and f_1 extends over the disks to give a homeomorphism. By induction this homeomorphism is isotopic to a composition of Dehn twists, so (after a bit of finagling to push the annuli of the Dehn twists off of the glued-in disks) the map f_1 on F_g is a composition of Dehn twists, so by composing these with the inverses of the D_k (which the reader can easily check are the Dehn twists along the same curves in the opposite direction), we get that f is isotopic to a composition of Dehn twists.

To prove the claim, though, we need to establish yet another claim first:

Sub-claim: Given any two non-separating simple closed curves γ, γ' in F_g meeting one another transversely (in a finite number of points), then there is a sequence of simple closed curves $\gamma = \gamma_1, \dots, \gamma_n = \gamma'$, each meeting its predecessor transversely in a single point.

Before proving it, let's see how it implies the claim. The basic idea is that if $|\gamma \cap \gamma'| = 1$, then we can take γ' to γ by a composition of two Dehn twists (and so the claim is proved by induction). To see this, look at a picture of a small neighborhood of the union of the two curves; it is a once-punctured torus. If orient the two curves, and then take one of the two curves and look at the image of it under the Dehn twist about the other (see the figures below), then it is easy to see that they are isotopic! Consequently, the composition of a Dehn twist about one with the Dehn twist about the other in the opposite direction take the other curve to the one curve! Composing the pairs of Dehn twists together for each adjacent pair therefore takes γ' to γ .





The sub-claim will be proved by induction! (Do you understand why everybody says that all proofs in surface topology are by induction?) We'll prove the inductive step first: it is stated as follows.

Inductive Step: Given two non-separating simple loops γ, γ' intersecting transversely in 2 or more points, then there is a non-separating simple loop γ'' such that

$$|\gamma \cap \gamma''|, |\gamma' \cap \gamma''| < |\gamma \cap \gamma'|.$$

This, together with the fact that given two disjoint non-separating loops, there is a third loop intersecting each of the first two exactly once (which we prove last), gives the sub-claim.

To prove the sub-claim, orient the two loops and start walking along one (γ , say) starting at some arbitrary point and look at the first two times you run into γ' . Up to reflecting the picture, the orientation of γ' gives rise to two cases: see the figure below.




In the first case we can simply take γ'' to be the 'dotted' curve shown in the figure. Because γ' is a connected loop, starting from one point of intersection it hits the other before it returns to the first, so γ'' hits γ at least one fewer time than γ' does, while (transverse) orientations (along γ') implies that γ'' hits γ' exactly once, verifying the inequalities above. Finally, since γ'' hits γ' once, it must be non-separating (otherwise, the first point of intersection represents γ' passing from one of the resulting two components to the other, and it can't get back!).

In the second case we must be slightly cleverer. Consider the two simple curves γ_1, γ_2 given by pinching γ' together (so the orientations match up). Each of them miss γ' completely (again by transverse orientation considerations), while each hits γ in at least

two fewer points. All we need, therefore, is that at least one is non-separating. For this we turn to first homology. $[\gamma'] \in H_1(M)$ is non-zero; it's not null-homologous since it is non-separating. And from the figure, it is easy to see (if you know what it means) that $[\gamma_1] + [\gamma_2] = [\gamma']$, since they cancel (homologically) along the arc between the two points of intersection of $\gamma \cap \gamma'$. So since the sum is non-zero, at least one of them is non-zero in homology, which means it is non-separating!

This leaves the case that γ and γ' are disjoint. But since both curves are non-separating (so the complement of each is connected), either the complement of both is connected, or the complement of both consists of two components (if you want to get really technical, this can be seen by using a Mayer-Vietoris sequence in homology for the pair $F_g \setminus \gamma, F_g \setminus \gamma'$). In the first case, pair off the resulting boundary components, one from each loop, and join them by an arc. If they intersect, by taking one 'around the bend' (see below) we can make them miss one another; then by gluing the two arcs together (after gluing the loops from γ , ~~same prop~~ back together), we get a loop hitting each exactly once. By shortcutting this loop (below), we get a simple loop with the same properties.


In the second case, it must be that the two components must contain one each of the loops from each splitting (the alternative is absurd: either each component has two loops, both from the same loop, and gluing back clearly gives a non-connected surface, or one component has one boundary component, clearly showing that the corresponding loop is separating). So we can draw the arcs as above in each connected component and glue together to give the required simple loop.



We have therefore proven the required claim and sub-claim of the sub-fact, so we have proven the sub-fact as well. Now, what does it have to do with links in S^3 and M ? Well, the map coh^{-1} can now be represented as a composition of Dehn twists, so we can change our picture of the gluing for M to the one below. Each of the Dehn twists D_k are supported on an annulus about curves γ_k (shown as a line segment and a point, respectively, in the figure). If we drill out a solid torus around each of the curves γ (which amounts to drill out a trough under each in the corresponding tops and bottoms of the small products), then the picture looks like the one on the right; the resulting gluings look as if they were by the identity.

(*) the resulting two arcs α, β , that is

But if we actually glued the products instead by the identity (the second set of figures), we get S^3 , and drilling out the solid tori again we get the picture on the right. But this is the same picture as the one above it! Consequently, M with the solid tori around the curves γ_k drilled out is the same as S^3 with the solid tori around the 'same' curves drilled out! This establishes the fact we have used.



So we have completed our proof (modulo the braids result) that all closed orientable 3-manifolds admit codimension-1 foliations. It is interesting to note that we could make the foliations we have built transversely orientable; verifying this observation is left to the reader. We could in fact spin the foliation (in our turbulizations) around each component of the links used in opposite directions, in various combinations, to create many different foliations. It is also easy to create many simple loops running transverse to the foliations we have built in M (so by turbulizing create many more foliations of M); start at a point and trace out a curve by following the leaf of the transverse 1-dimensional foliation that we are provided with - either we get really lucky and the curve closes up of its own accord, giving us a transverse loop, or it doesn't. But then it must pass arbitrarily close to itself (since we are in a compact manifold), so we can 'short-circuit' it (like we did before in proving that spaces of leaves in \mathbf{R}^2 are manifolds) to create a transverse loop. (This argument uses the fact that we know we have a transversely-orientable foliation.) these hypothetical transverse loops could hit any of the torus leaves of the Reeb components, a fact that will become very relevant later on.

Foliations and the Topology of 3-manifolds

Outline of class 7

Today we get to the real core results of this class. We saw last time that every closed, orientable 3-manifold admits a (transversely-orientable, in fact) codimension-1 foliation. (It's actually true that closed, non-orientable 3-manifolds admit codimension-1 foliations (there are two proofs, one by Woods and the other by Thurston), but we won't present them here.) From the point of view, however, that we will be taking in this class, this result that we have just proved is an out-and-out tragedy. We want to use foliations to tell us something about what a 3-manifold looks like, and this result tells us that anything we could prove would have to be true of all closed 3-manifolds M . And about the only thing that fits that bill is that $\chi(M)=0$ (which a foliation does tell us, but so what).

What we need to do now is to find some 'reasonable' extra condition or conditions to impose on a foliation to restrict the kinds of manifolds we can foliate in that or those ways. The hope is that by choosing the right sort of condition we will be able to prove that all of the 3-manifolds we get have some 'desirable' properties, from the 3-manifold topologist's point of view. What are some such properties?

1. $\pi_1(M)$ is infinite. This is a desirable property because it avoids all of the known counterexamples - Lens spaces - to a long-standing conjecture that homotopy-equivalent 3-manifolds are homeomorphic.

2. $\pi_2(M)=0$. This is desirable since together with (1) a quick exercise in topology and homotopy theory implies that the universal cover of M is contractible. Two such manifolds therefore are homotopy equivalent iff their fundamental groups are isomorphic.

3. M is irreducible. This means every embedded 2-sphere bounds an embedded 3-ball. This is actually (formally) stronger than (2); that fact is known as the Sphere Theorem. (Many people use that theorem (1957?) to mark the beginning of modern 3-manifold topology.) The converse to this theorem is basically the Poincaré conjecture.

4. The universal cover of M , \tilde{M} , is homeomorphic to \mathbf{R}^3 . It is actually a conjecture that (1) and (2) together imply (4) (if you want to avoid trying to prove the Poincaré conjecture in the process you can substitute (3) for (2)). A different way of stating it is that the only contractible 3-manifold that covers a compact 3-manifold is \mathbf{R}^3 .

These are all considered useful assumptions to make/goals to achieve when studying a 3-manifold. The amazing fact is that there is exactly one, easy to state, criterion that we can impose on a foliation to insure that our 3-manifold has all of these properties. To figure out what it is, all you need to do is go back and look at our previous construction, and see what made it so easy to foliate all of those 3-manifolds. What do all of the foliations have in common? The answer is that they all have Reeb components - a torus leaf, bounding

a solid torus, which is foliated with a Reeb foliation. Avoiding them is all it takes to get all of the properties above. These results are found in the following three theorems. We will assume (mostly for convenience) that all of the manifolds mentioned are orientable and the foliations are transversely orientable. (Removing these properties adds a few more manifolds to the ‘exceptions’ list.)

Theorem (Novikov): If M admits a foliation \mathcal{F} without Reeb components, then

1. $\pi_1(M)$ is infinite,
2. $\pi_2(M)=0$ (except for $M \cong S^2 \times S^1$), and
3. for every leaf L of \mathcal{F} , the inclusion-induced homomorphism $\pi_1(L) \rightarrow \pi_1(M)$ is injective.

Theorem (Rosenberg): If M admits a foliation \mathcal{F} without Reeb components, then M is irreducible (unless $M \cong S^2 \times S^1$).

Let’s look at property (3) of Novikov’s theorem. If we were to lift the foliation \mathcal{F} to a foliation $\tilde{\mathcal{F}}$ of the universal cover \tilde{M} , then the commutativity of the diagram below easily shows that any leaf \tilde{L} of $\tilde{\mathcal{F}}$ π_1 -injects into \tilde{M} , i.e., it is simply-connected. But the only simply-connected surfaces are S^2 and \mathbf{R}^2 . We will see that if $\tilde{\mathcal{F}}$ contains an S^2 -leaf, then (so does \mathcal{F} and) $M \cong S^2 \times S^1$ (which is exactly where our exceptional case comes from!), so we can assume in all other cases that every leaf of $\tilde{\mathcal{F}}$ is a plane. Then the following result becomes very interesting.

Theorem (Palmeira): If N^n is a simply-connected (open) manifold foliated by hyperplanes \mathbf{R}^{n-1} , then $N \cong \mathbf{R}^n$.

Setting $N = \tilde{M}$ from our discussion above, we immediately get


Corollary: If M admits a foliation without Reeb components, then either $M \cong S^2 \times S^1$ or $\tilde{M} \cong \mathbf{R}^3$.

We will start with Novikov’s theorem. We will actually prove its contrapositive:


Novikov’s Theorem: If M is an orientable 3-manifold with a transversely orientable foliation \mathcal{F} , and if

1. $\pi_1(M)$ is finite,
 2. $\pi_2(M) \neq 0$, or
 3. for some leaf L of \mathcal{F} , the inclusion-induced homomorphism is not injective,
- then \mathcal{F} contains a Reeb component (or in case 2., $M \cong S^2 \times S^1$).

We will prove this by developing a homotopy-theoretic criterion which implies the existence of a Reeb component, and then show that each of the homotopy-theoretic criteria above imply this new criterion. To see what this criterion will be, let’s take a look at a Reeb solid torus.



We can apply this reasoning, for example, to loops γ in leaves (which are formally maps $f:[0,1] \rightarrow L$ with $f(0)=f(1)$). Over (the image of) this loop we can erect a positive and negative normal fence by taking the union of the integral curves of the transverse vector field passing through each point of the image (see figure). Then lifting the loop to nearby leaves gives us a path in the leaf, lying in this normal fence. It will actually look like the intersection of this leaf with the normal fence (or the pullback, if you think of the fence as a map of an annulus into M) making one circuit around the fence. It then makes sense to talk about a return map associated to this loop; by walking around these lifted paths, we can define a map which assigns to the initial point of a path its ending point. This defines an injective map from a small neighborhood of 0 (thinking of the transverse curve as $[-1,1]$) to some other small neighborhood of 0. This is the holonomy map; it describes how the foliation changes, transversely, as we walk around the loop (that is what holonomy is; it describes how some quantity or collection of objects change as you walk around a loop).



For our purposes, we need only make two distinctions in the type of behaviors we get walking around loops. If, as we walk around a loop, in the positive or negative direction (as determined by the transverse vector field) all lifted paths sufficiently close to γ close up to form loops in their leaves, we say that γ is a (positive or negative) **non-limit cycle**. If in the positive or negative direction no matter how close to γ we look there are always lifted paths which do not close up, we say that γ is a (positive or negative) **limit cycle** (see figure). (In more modern terminology we would say the γ has trivial holonomy (on the positive or negative side), and that γ has non-trivial holonomy). It's not hard to see that the property of being a non-limit cycle is invariant under taking free homotopies in the leaf L , and that the actual holonomy around a loop is invariant under based homotopy in the leaf - it's just a matter of using the 'right' map of a 2-disk (as below), and interpreting the results the right way; we leave it as an exercise (since we won't really use either fact).

It's easy to see that the existence of a Reeb component in fact implies the third statement of the theorem: The torus leaf T fails to π_1 -inject. In fact we can see a disk in the solid torus (the meridian disk) which meets the only its boundary, and that boundary loop γ is not null-homotopic in T , but the disk exhibits the fact that it is null-homotopic in the solid torus (and hence in the 3-manifold M that the solid torus sits in). This **compressing** disk meets the other \mathbf{R}^2 -leaves in concentric loops parallel to γ , and, since these loops are in simply-connected leaves, they are in fact null-homotopic in their leaves. This fact is the key to our homotopy-theoretic criterion. But to describe it more exactly, we need to introduce the notion of holonomy.

Holonomy.

The basic idea behind holonomy is that if we pick a leaf L of a codimension-1 foliation \mathcal{F} of M , then the other leaves of \mathcal{F} that pass near L look, in small pieces, like they are covering spaces of L (by projecting down along the 1-dimensional foliation transverse to \mathcal{F} - this is a local diffeomorphism). Therefore things like lifting objects and maps in L to nearby leaves should act a lot like the lifting criterion for covering spaces. We'll now proceed to make this idea more rigorous.

We will deal almost exclusively with 'lifting' maps $f: A \rightarrow L$ of contractible sets (like arcs and disks) to a leaf up to nearby leaves. Here the idea in covering space theory is that you picture your favorite basepoint and lift it, and then decide how to lift another point by drawing a path to it from the basepoint, and deciding how to lift it by dragging yourself along the path - the local triviality of the projection ensures that you never have to make any decisions how to proceed. The same idea works here, because the local coordinate charts ensure that, in the small, there are no decisions to make on how to lift the map. You simply choose basepoint for A and lift its image to a nearby leaf, and then to lift any other point you draw a path to it and walk along the path, piecemeal lifting it to the nearby leaf by using the fact that in the small, our foliation just looks like horizontal planes, so it's easy to tell how to extend. The fact that this lifting doesn't depend on the path chosen is because any two are homotopic rel endpoints, which (as in an earlier construction) we can realize as a sequence of 'small' homotopies (the images of whose supports are each contained in a distinguished neighborhood), so it is easy to see that the endpoint of the path (i.e., which plaque we end up in) hasn't changed.

In fact the only thing we need to worry about is that the leaf we are lifting to as we drag along a path and the leaf we started in continue to hit the same distinguished charts (see figure). This of course need not be the case; but by only trying to lift maps to leaves very close to L (where close must be interpreted relative to how complicated the map f is, i.e., how far across L the map wanders), we can always insure that we pass from one 'good' chart to another', so can always lift entire paths.



One easy way to insure that we have a non-limit cycle is to look at a loop γ null-homotopic in the leaf L . Then there is a map $f:D \rightarrow L$ of a 2-disk into L with $f|_{\partial D} = \gamma$. By our 'lifting criterion', this map lifts to a map of a disk in all nearby leaves; in particular, the lifted map restricted to the boundary is a lift of γ to a closed loop. Also, these lifted loops are also null-homotopic in their leaves (since they extend to the lifted maps of the disk).

A vanishing cycle is basically like a null-homotopic loop, in that it lifts to null-homotopic loops, it's just not itself null-homotopic.

Definition: A **vanishing cycle** for the foliation \mathcal{F} of M is a (singular) loop γ in a leaf L such that γ is a (positive or negative) non-limit cycle, the lifts of γ in the appropriate direction are all loops null-homotopic in their respective leaves, but γ is not null-homotopic in its leaf L .

Note that this is exactly true of the meridional loop in the torus boundary of a Reeb component. Note also that the definition only talks about null-homotopic versus homotopically essential loops - it's just a homotopy-theoretic criterion on our leaves.

The proof of Novikov's theorem naturally breaks into three pieces:

1. Show that each of the three conditions in the theorem implies the existence of a singular (although by general nonsense we can assume it is immersed) vanishing cycle in some leaf of \mathcal{F} (this basically turns one homotopy-theoretic fact into another),
2. Show that if some leaf has an immersed vanishing cycle, then some leaf has an embedded vanishing cycle (this is the heart - it turns a homotopy-theoretic fact into a geometric one), and
3. Show that a leaf with an embedded vanishing cycle is a torus, bounding a solid torus with a Reeb foliation (which turns one geometric fact into another).

We'll handle these in the order 3., then 1., then 2. (so that we get a chance to warm up a bit). Actually, 1. is not quite exact - we'll show that under the second condition ($\pi_2(M) \neq 0$), either \mathcal{F} has a vanishing cycle or some leaf of \mathcal{F} has non-trivial second homotopy group. Now the only surfaces with this property are S^2 and \mathbf{RP}^2 (and under our orientability hypotheses, an \mathbf{RP}^2 leaf cannot occur). Since our exceptional case stems from this one possibility, that some leaf of the foliation is a 2-sphere, we will deal with this exception first (and, as it happens, develop some useful techniques for dealing with part 3. of the theorem). The result we will prove is called the Reeb Stability Theorem. We'll begin with it next time.

Foliations and the Topology of 3-manifolds

Outline of class 8

Today we will prove the following theorem:

Theorem (Reeb Stability): If a closed orientable 3-manifold M admits a transversely-orientable foliation \mathcal{F} , and one of the leaves L of \mathcal{F} is a 2-sphere, then $M \cong S^2 \times S^1$, every leaf of \mathcal{F} is a 2-sphere, and the foliation consists (up to homeomorphism) of leaves $S^2 \times \{\text{pt.}\}$.

We can relax the orientability and transverse-orientability assumptions without too much difficulty (by passing, if necessary, to a 2- or 4-fold covering); we then get that any closed 3-manifold admitting a foliation with an S^2 or \mathbf{RP}^2 leaf is covered, in a leaf-preserving way, by $S^2 \times S^1$; it is an easy matter then to determine what possible manifolds there are (using a little bundle technology). They are the two bundles over S^1 with fiber S^2 , the two bundles over S^1 with fiber \mathbf{RP}^2 , and the connected sum of two \mathbf{RP}^3 's ($\mathbf{RP}^3 \setminus (3\text{-ball})$ can be foliated with one \mathbf{RP}^2 leaf, the rest being 2-spheres).

We will see that not only will this result be useful to us (it describes exactly where the exceptional cases of our main theorems come from), but the techniques we develop during its proof will carry over almost unchanged to the last third of the proof of Novikov's theorem.


We will prove Reeb Stability by showing that the union \mathcal{U} of the set of 2-sphere leaves of \mathcal{F} is both open and closed in M (hence is all of M , since M is (assumed) connected, so every leaf is a 2-sphere), and the space of leaves is a (compact Hausdorff) 1-manifold, hence a circle. From there it is easy to see that M is one of the two bundles over S^1 with fiber S^2 , so, since it is orientable, $M \cong S^2 \times S^1$.

The first part, that \mathcal{U} is open, is not hard, and in fact follows immediately from holonomy considerations. Pick a 2-sphere leaf S , and write it as $S = D^+ \cup D^-$, the union of its northern and southern hemispheres. These are embedded disks in S , so they lift along the transverse vector field, in both directions, to embedded disks in nearby leaves, D_ϵ^+ , D_ϵ^- . Consequently, their common boundary, $D^+ \cap D^- = \gamma$ lifts to loops in nearby leaves γ_ϵ , which is the common boundary of the lifted disks. Therefore, every nearby leaf is the union of two disks glued together along their boundaries, i.e., is a 2-sphere. So every leaf sufficiently close to a 2-sphere leaf is a 2-sphere; the set \mathcal{U} of 2-sphere leaves is open.

Now suppose that \mathcal{U} is not closed; so there is a point $x \notin \mathcal{U}$ which is a limit point of \mathcal{U} . In particular, the leaf L of \mathcal{F} containing x is not a 2-sphere. Therefore there is a sequence of points x_n in \mathcal{U} limiting down on x . In particular, in some distinguished chart about x , the plaques containing the x_n must be limiting down on the plaque containing x (they can't be in the same plaque as x , since then the plaque containing x would be in a 2-sphere leaf!). We can therefore without any loss of generality replace the x_n with the points in

their plaques lying directly above x (in the direction of the transverse vector field); this sequence (is also in \mathcal{U} and) also limits on x .



We can therefore imagine ourselves to be in the situation pictured below. As we stare down the transverse foliation, we can be standing on x_n , a point in a 2-sphere leaf, but a short distance away, the point x is not in a 2-sphere leaf. Since \mathcal{U} is open, it intersects the short arc between x_n and x in an open set, so there is a first point z along this arc going from x_n to x which is not in a 2-sphere leaf (possibly $z=x$). We will demonstrate that z is in fact in a 2-sphere leaf, giving us a contradiction to the assumption that \mathcal{U} is not closed.




We will do this by constructing a map $f: S^2 \times [0, \epsilon] \rightarrow M$ which maps 2-sphere levels injectively into leaves of \mathcal{F} , for which $z \in f(S^2 \times \{0\})$. Since this image is the continuous image of a compact set into a Hausdorff set (the leaf L) under an injective map, it is a subset of L homeomorphic to S^2 ; since L is connected, we therefore have that L is a 2-sphere.

The idea is to just start at the 2-sphere leaf S^2_ϵ containing x_n and flow along the transverse foliation in the direction of z , parametrizing the resulting map by what leaf we land in (using a parametrization of the short arc from x_n to z (letting z correspond to 0)). The point is that unless as we flow along the trajectories from a point of S^2_ϵ the path we trace out gets infinitely long before we reach zero, we will be able to define our function for all $t > 0$, as well as define a limiting value for $t = 0$. (This is because we can always push a little bit further from any point we can reach - we are just flowing along the transverse vector field. There is nothing that we can run into to stop us.) It will then be a relatively easy matter to show that all of the properties claimed for the resulting function are true.

But there is the possibility that one of the paths traced out in this way reaches infinite length before it reached zero (because we are imposing an unnatural parametrization on the path - it is parametrized by the leaves it passes through). What we will show is that if this were to happen, then we would have already passed through every leaf of \mathcal{F} . So they are all 2-spheres - they are all among the 2-sphere leaves lying between x_n and z !



We've already used a variant of the argument required for this as well. The idea is that if our trajectory has become infinitely long, then since we are in a compact manifold M , this path α must have passed arbitrarily closed to itself; in particular, it must have passed through the same distinguished chart twice (see figure above). Therefore it must have passed through the same 2-sphere leaf $S_r^2 = S_s^2$ twice. But if we then take the set \mathcal{V} of 2-sphere leaves lying between the points corresponding to r and s in our parametrizing arc, we can easily see that this set is both open and closed in M (hence is all of M). It is closed because for any sequence in \mathcal{V} converging in M , $a_n \rightarrow a$, using the same idea as above, we can assume that the a_n are converging straight down along the transverse foliation. But then the trajectory they are following reaches a without becoming infinitely long, so the parameters corresponding to the a_n (which lie in $[r, s]$) must converge to that of a , so it also lies in $[r, s]$, i.e., $a \in \mathcal{V}$. On the other hand, the set is open, since if we pick any point in S_t^2 for $t \in (r, s)$, we can draw a path in its leaf back to the point t in the parametrizing arc. There an open set around the point is obvious, and it can be dragged back along the path to give an open neighborhood of our point (see figure below). If on the other hand the point lies in $S_r^2 = S_s^2$, then we can draw two paths, one each back to r and s . Then we can take two half-neighborhoods (one above and one below) and drag them back along the paths to give two half-neighborhoods (one above, one below) around our point; inside their union it is easy to see that we can find an open neighborhood of the point (see figure below).




So we can assume that all of our trajectories from S_ϵ^2 are defined for all $t \geq 0$, and have finite length. But then it is easy to see that the map $f: S_\epsilon^2 \times [0, \epsilon]$ given by $f(x, t) =$ the point along the trajectory from x in the leaf S_t^2 is a continuous function (this follows readily from a distinguished coordinate picture, including the limiting case $t=0$, since we can imagine the transverse foliation consisting of vertical lines (or at least as continuous paths transverse to the horizontal plaques); the parametrization of the horizontal plaques by which 2-sphere they are in is a strictly monotone decreasing continuous function). Notice that we can assume that each of our 2-sphere leaves corresponds to a unique parameter - otherwise, a trajectory passes through a leaf twice, and the argument above applies instead, to get what we want!

Therefore the restriction of this function to the sphere $S_\epsilon^2 \times 0$ maps it continuously into the leaf L containing z . This map is injective (by the uniqueness of solutions of ordinary differential equations, basically); no two distinct trajectories pass through the same point. The only possibility then is that we limited on this point from both sides (see figure below);

but this violates our transverse orientability assumption (this is actually the only place where we use it!) - the non-zero vector at that point would have to be pointing both ways.

We therefore have verified all of the properties used in our argument above, so the leaf L containing z is a 2-sphere, a contradiction. So every leaf is a 2-sphere.

Finally, showing that the space of leaves is a Hausdorff manifold is easy; it is locally Euclidean since each 2-sphere leaf pushes off of itself to nearby 2-spheres, so a neighborhood of every leaf looks like S^2 crossed with an open interval. The transverse interval projects injectively to an open neighborhood of the point corresponding to our leaf; it is an open set since its inverse image is our neighborhood of the leaf, so is a saturated open set. Finally the space of leaves is Hausdorff since all of our leaves are compact; any two distinct leaves have disjoint ϵ -neighborhoods, and each contains a saturated neighborhood like we just described (details left as an exercise - each leaf is covered by finitely-many distinguished charts). Finally, the space of leaves is compact, since it is the continuous image of M under the quotient map.



That finishes the proof of Reeb stability. It is actually true more generally that for any codimension-1 foliation \mathcal{F} of a compact manifold the union of the compact leaves of \mathcal{F} forms a closed set, and that in fact for any compact leaf any other compact leaf sufficiently close to it is diffeomorphic to it (so as a consequence any foliation has only finitely many compact leaves, up to diffeomorphism, and the union of the set of compact leaves diffeomorphic to a given one is also closed (which is what we proved for S^2 above). The proof is a variant of what we gave above, but is a bit more involved. Also, the openness part of our prove can be generalized (this generalization is actually what is usually called Reeb stability) - if a compact leaf contains no loops with non-trivial holonomy, then every leaf sufficiently close to it is diffeomorphic to it. This can be strengthened somewhat to include the case that the set of leaves with non-trivial holonomy forms a finite subgroup of the fundamental group of the leaf. We will probably not pursue these generalizations.

Our next task will be to utilize this idea of flowing along the transverse foliation to prove the last third of our outline of Novikov's theorem. In that case however, we will find that some trajectory has to become infinitely long before reaching 0; it will be an essential ingredient to the proof! compact

Foliations and the Topology of 3-manifolds

Outline of class 9

Today we will prove the third step in our program to prove Novikov's theorem. We will show that if a leaf L of a foliation \mathcal{F} of M contains an embedded vanishing cycle γ , then L is a torus bounding a solid torus with a Reeb foliation. The techniques involved are similar in spirit to our proof of Reeb stability, but, as it turns out, the conclusions are exactly the opposite.

If we look at the normal fence over $\gamma=\gamma_0$, nearby leaves, by definition, meet it in loops γ_t which are null-homotopic in their leaves. If we choose a short enough fence, we can assume that its image is an embedded annulus A (since γ is embedded). Now each γ_t for $t>0$ is an embedded null-homotopic loop in the leaf L_t ; consequently, each bounds a disk D_t in the leaf.

One way to see this: lift the loop up to the universal cover of the leaf; it is \mathbf{R}^2 , so the Jordan curve theorem implies that it bounds a disk. This is true for every lift, and further, the disks bounded by each lift are disjoint. For otherwise, since their boundaries are disjoint (so one is contained in the other), and since there is a covering translation taking any lift to any other, this translation would carry the one disk into itself. Therefore the covering translation has a fixed point by the Brouwer Fixed Point Theorem, a contradiction (the only one which fixes a point is the identity!). From here it is easy to see that the covering projection maps any one of these disks injectively down to L_t , giving a disk bounded by γ_t (otherwise, there is a point downstairs covered by two upstairs - there is a covering translation carrying one to the other, which, like before, must have a fixed point).

Now pick one of the disks D_ϵ , and look at $D_\epsilon \cap A \subseteq D_\epsilon$. This is a collection of loops (actually they are all γ_t 's), in a disk, so we can pick an innermost one γ_{t_0} . Then the union of the disk (in D_ϵ , actually) and the annulus in A between γ_{t_0} and γ_0 embedded in M (since each piece is embedded and meet only in γ_{t_0}). Now γ_{t_0}, ϵ . ^{rename} _{is}

Now start flowing from D_ϵ down toward the leaf L_0 containing the vanishing cycle γ_0 , as we did in proving Reeb stability. We can then build, as we did then, a map $f: D_\epsilon \times (0, \epsilon] \rightarrow M$, and try to extend it to 0. The map obtained by flowing along trajectories is defined for all positive t , since the set of t for which the function can be defined for all points of D_ϵ is open (because any disk D_t we can get to we can use holonomy to flow off a little further); but if there is a first time $t_0 > 0$ which we can't reach, ^{then} since we can flow both backwards and forward from the resulting disk D_{t_0} , we can go back slightly to a point we have defined the function for, and then flow forward past the point we thought we had to stop. α^t .

But unlike the case of Reeb stability, we can't extend the map to $t=0$ for all points of D_ϵ ; otherwise, as in the proof of Reeb stability, this would give a continuous map of a disk


into L_0 with boundary γ_0 , proving that γ_0 is null-homotopic in its leaf, a contradiction. This means, therefore, that there has to be a point $x \in \text{int}(D_\epsilon)$ whose trajectory becomes infinitely-long (just) before we reach $t=0$. This trajectory, since M is compact, must be piling up somewhere - it passes through the same distinguished chart infinitely-often. It therefore in particular passes through the same plaque infinitely-often, and so there is a sequence $f(x, t_n)$ limiting on a point y , for a sequence $t_n \rightarrow 0$.

But this plaque sits on the leaves L_{t_n} , and in fact we must eventually have $y \in D_{t_n}$, since otherwise the loops γ_{t_n} must always lie between $f(x, t_n)$ and y , so there is a sequence of points in γ_{t_n} limiting on y , which is impossible (since these loops are a bounded ^(below) distance apart (they lie above one another in the transverse direction)). But then one of the trajectories from D_ϵ must pass through y (since this is true of every point of the D_t 's), and we can then follow it back to D_ϵ . But then every trajectory passing sufficiently close to y (in particular, the one we have been looking at!), flows back (over a finite distance, so has only bounded spread from the one through y) to hit D_ϵ in the same amount of 'time'. Therefore the trajectory through x (since it passes arbitrarily close to y infinitely-often) hits D_ϵ infinitely-often. Therefore there is a first time t_0 in which it happens.

We therefore have for this time t_0 that $D_{t_0} \cap D_\epsilon \neq \emptyset$, and so since $\gamma_\epsilon \cap \gamma_{t_0} = \emptyset$ (they lie in different levels of the annulus, which is embedded), one of these disks contains the other. But if D_ϵ contains D_{t_0} , then it also contains γ_{t_0} , implying the $D_\epsilon \cup A$ is not embedded; so we have that $D_\epsilon \subseteq D_{t_0}$.

Exercise: Show that the first time that a trajectory hits D_ϵ again is the same for all such trajectories (Hint: look at the argument below, and notice that it shows that any trajectory, which misses D_ϵ when another one hits, never hits D_ϵ (there is a first time one of them hits D_ϵ again)).

Therefore, if we imagine the image $f(D_\epsilon \times [t_0, \epsilon])$, the top $D_\epsilon \times \{\epsilon\}$ gets glued to a disk in the interior of the bottom. It is therefore a solid torus 'bent' at the annulus $f(\partial D_\epsilon \times [t_0, \epsilon]) \simeq A$ (see figure below).




But if we look at the other annulus $B = D_{t_0} \setminus \text{int}(D_\epsilon)$, every trajectory passing through

those points must have a limiting value (i.e. they have finite length). This is because otherwise, as above, each trajectory would have to pass through the disk D_ϵ infinitely often. But up until it hits the annulus B it could have hit only finitely-many times (since the trajectory has finite length up until then) and the transverse vector field points outward all along the boundary of the solid torus (except along A where it is parallel), so once you go out of the solid torus (which is where the trajectories are going) they can never get back in. So they clearly can't hit D_ϵ infinitely-often.

But now if we look at the image of the annulus B as $t \rightarrow 0$, we get an annulus in L_0 , except that its ends (both lie on A and therefore) get glued to one another (see figure below). Since the points on B are all distinct, their limits are all distinct, so the only identifications occur at the ends, so we end up with an (embedded) torus contained in L_0 ; consequently L_0 is a torus.

So we have a solid torus $f(D_\epsilon \times [t_0, \epsilon])$ surrounded by a torus L_0 ; in between we see the annulus B crossed with an interval, glued together along $\partial B \times I$ (with a shift designed to avoid the corner in the solid torus). Therefore, the union of the two pieces is a solid torus bounded by the torus L_0 . The leaves in the solid torus consist each of one of the disks D_t , $(t_0 < t \leq \epsilon)$ which have boundary in the corner of the interior solid torus, glued to a sequence of annuli (flowed off of the annulus B) which string together to spiral out to the torus L_0 (see figure). In other words, this is a Reeb foliation of the solid torus!



This completes the proof of the third segment of our outline of Novikov's theorem. We'll finish with a few observations about the Reeb foliation.

Fact: If a foliation \mathcal{F} of M has a Reeb component bounded by the torus leaf T , then no loop in M everywhere transverse to \mathcal{F} passes through T .

We more or less already noticed this before: we can assign a normal orientation to the leaves in the solid torus bounded by T so that it points everywhere inward along T (even though the orientation might not extend over the entire transverse foliation in M). Any tangent vectors to an (oriented) transverse arc which passes through T (imagine it

passing into the solid torus) starts and stays within 90 degrees of the transverse vector field (otherwise it must pass through 90 degrees, hence is tangent to the foliation). But in order to get out of the solid torus again, it would have to make an angle of greater than 90 degrees with the normal orientation, an impossibility. So once you get into the solid torus, you can't get out again; so an transverse path passing through T can't be completed to a transverse loop.

This fact prompts a definition:

Definition: A codimension-1 foliation \mathcal{F} of M is called **taut** if for every leaf L of \mathcal{F} there is a loop γ transverse to \mathcal{F} which passes through L .

The above argument shows that taut foliations have no Reeb components, and therefore enjoy all of the properties which we are proving about foliations without Reeb components. In fact, we will, after proving those results, begin to focus our attention almost entirely on taut foliations. They are the kinds of foliations which Dave Gabai developed powerful techniques for constructing, and used to study knots in the 3-sphere (in so doing, settling several long-standing conjectures). It is the techniques that he (and Thurston) developed that will occupy us in the second half of this class.

We should notice in closing that the statement above is not reversible; there are foliations which are not taut, and yet contain no Reeb components. The easiest way to do this is to build a transversely-oriented foliation with a separating compact leaf. For example, take a once-punctured torus T_0 , and foliate two copies of $T_0 \times S^1$ by $T_0 \times \{\text{pt.}\}$, and then 'spin' each foliation at the boundary, to make the boundary a leaf (we could think of this as taking $T^2 \times S^1$ with foliation by $T^2 \times \{\text{pt.}\}$, and turbulize along a loop $\{\text{pt.}\}$ and throw away the Reeb solid torus). Then glue the two foliated manifolds together.

This is in some sense the only way for the reverse to fail, in fact:

Fact (Goodman): Any leaf (in a compact manifold) which has no transverse loop passing through it is a torus.

This to me is an amazing fact; it says for example that there is no way foliate a genus- k handlebody (for $k > 1$) with the boundary a leaf (otherwise (after passing to a suitable covering to get a transversely-orientable foliation) you could glue two together to get a foliation with F_k as separating leaf). Notice that the $k=1$ case is an exception - it can be given a Reeb foliation!

Exercise: Show that any non-compact leaf in a compact manifold has a transverse loop passing through it (Hint: it passes through a distinguished chart twice; pass to a transversely-orientable foliation first). Can you find a non-taut foliation with only non-separating tori? Think of ('Reeb') foliations of annuli, and build one in the 3-torus.

The proof of Goodman's result is analytical (you integrate a cohomology class along the leaf to see that the Euler characteristic is 0); we'll bump into some of the pieces later on.

Foliations and the Topology of 3-manifolds

Outline of class 10

Last time we saw how our proof of Novikov's theorem will end: the existence of an embedded vanishing cycle implies the existence of a Reeb component. Today we will begin to see how our proof will begin: how each of the three conditions of Novikov's theorem (almost) implies the existence of an (immersed) vanishing cycle.

For this we will need to introduce a new technique: the notion of the pullback (singular) foliation. The basic idea for this is that codimension-1 foliations are the natural setting for something like a Morse theory of functions $f:F^2 \rightarrow M^3$ to make sense.



A **Morse function** is a smooth map $f:F^2 \rightarrow \mathbf{R}^3$ from a surface F to Euclidean space so that the critical points of $p \circ f : F \rightarrow \mathbf{R}^1$ (where p is the projection onto the third coordinate) are discrete (so, if F is compact, finite) and **non-degenerate**. For our purposes, this means that in a neighborhood of a critical point, the map looks like one of the following pictures:



(Formally, these are supposed to be quadratic surfaces; paraboloids and hyperboloids. We will only be concerned with their qualitative shape, however.) The idea is that we are in some sense trying to make the map f transverse to the horizontal foliation (meaning that $f_*(T_x F) + T_{f(x)} \mathcal{F} = T_{f(x)} \mathbf{R}^3$). The point is that this fails ($f_*(T_x F) = T_{f(x)} \mathcal{F}$) only finitely-often, and at the failures you get one of our pictures.

The main technical facts we need about such maps (and which we will not prove) is that they form an open and dense subset of the set of all continuous maps from F to \mathbf{R}^3 . So any map can be deformed (through arbitrarily small deformations) to a Morse function, and small enough deformations of Morse functions remain Morse. Let us now apply these facts to the more general problem of making a map $f:F \rightarrow M$ transverse to a foliation \mathcal{F} of M .

Since F is (assumed) compact, its image in M is compact, so we can cover the image by finitely-many distinguished charts. Then F can be cut up into a finite number of small enough pieces so that each piece maps into one of the distinguished charts. So each piece can be thought of instead as being mapped by f into \mathbf{R}^3 , where we have a horizontal foliation. By a small deformation of f on one of the pieces (which we extend to a small deformation across F), we can make f restricted to this piece a Morse function into the distinguished neighborhood. By doing this to each piece in turn (by smaller and smaller deformations, so that the map on the pieces we have already dealt with stays Morse) we can then by induction deform the map f so that on every piece, the map to the corresponding distinguished neighborhood is Morse. We can then use this function to **pull back** the foliation \mathcal{F} on M to a **singular** foliation on F . We do this by doing it locally; on each piece, the map to the corresponding chart is transverse to the foliation by horizontal planes (i.e., to \mathcal{F}) except at a finite number of points. Therefore, except at the critical points, the projected map to \mathbf{R}^1 is a submersion, so the domain can be foliated by level sets $(p \circ f)^{-1}(*)$ (which actually is the same as $f^{-1}(\text{horizontal plane})$), at least, off of the critical points. These foliations on the pieces fit together to foliate F off of the critical points, because the transition functions of the coordinate charts take horizontal planes to horizontal planes, so they preserve the level sets of their corresponding $p \circ f$'s. If we add the critical points, we get a singular foliation of F rather than a foliation; but by understanding what f looks like at the critical points, it is easy to see how the level sets nearby are behaving, so we have a very good picture of the foliation near the 'center' and 'saddle' singularities (see below).



So we get this singular foliation of our surface F . We call it the pullback foliation of \mathcal{F} under f , and denote it $f^*(\mathcal{F})$. One last technical point that we need: we can arrange the map f (by a further small deformation) so that no two saddle singularities are joined by an arc of the singular foliation. We do this by making sure that none of the critical points of our maps on pieces lie in the same leaf of \mathcal{F} . If two do, then since a single leaf can hit a chart in only countably-many plaques, there are other leaves arbitrarily close to each of the critical points, and by 'lifting' on or the other up (by carrying out a small deformation supported on a neighborhood of the critical point; see figure below) we can push the critical point into a different leaf, making it impossible to be joined by an arc

(the arc maps into a single leaf). We can, in fact, just to be on the safe side, deform the map f so that all critical points are in distinct leaves. Since this can be done for free, we assume we have done so.



The whole point to this kind of pullback foliation is that we can now use it to carry out an Euler characteristic calculation of the surface F . Since our foliation F is (assumed) transversely-orientable, we can use the transverse orientation of \mathcal{F} to transversely-orient the pullback away from the singularities (choose the normal direction which is mapped under f to a vector pointing to the same side of the leaf we're standing on as the transverse orientation of \mathcal{F}). Then since F is (assumed) orientable, we can use this transverse orientation as the first half of an orientation frame each point of F (away from the singularities), then using the second half of the frame to tell us which way to turn, we can rotate our normal orientation to orient the leaves of our singular foliation on F (again, away from the singularities). In other words, we can think of our pullback foliation as having come from a vector field (with zeros - the singularities are the zeros of the corresponding vector field). But a vector field with isolated zeros (such as we have) can be used to calculate the Euler characteristic of our surface F .

A quick look at the neighborhoods of our two types of singularities makes it evident what the index of these zeros of the vector field are. At the center singularities we have an index 1 zero, at the saddle singularities we have an index -1 zero. Therefore, the Euler characteristic is the sum of these 1's and -1 's as we range over the singularities. But we actually use this calculation in reverse; we will know what the Euler characteristic of our surface is; this calculation tells us what kinds of singularities we must have! In particular, since for our proof of Novikov's theorem we will be applying this to (2-spheres and 2-disks doubled along their boundaries, which are) 2-spheres, which have Euler characteristic 2, and since you can't get 2 by just adding up a bunch of 1's, we can conclude that every (Morse) singular foliation of the 2-sphere has at least 2 center singularities.

The way we will use this fact is that each of the three conditions of Novikov's theorem gives us a map of a 2-disk or a 2-sphere, where each satisfies its own special group-theoretic

condition. In the first, $\pi_1(M)$ finite, since we can always find a loop transverse to our foliation, some power γ of it is null-homotopic, so there is a map $f:D_2 \rightarrow M$ such that $f|_{\partial D} = \gamma$ is transverse to the foliation. In the second case, we get a map of a 2-sphere into M which is not null-homotopic. If we assume that M is not $S^2 \times S^1$, then (by Reeb stability) none of the leaves of \mathcal{F} are S^2 's or \mathbf{RP}^2 's, so every leaf L (has universal covering \mathbf{R}^2 , and so) has $\pi_2(L)=0$; so in particular, the map M cannot be deformed to a map into a leaf (otherwise it could then be deformed to the constant map inside that leaf). In the third case, we get a map $f:D_2 \rightarrow M$ such that $f|_{\partial D} = \gamma$ is a loop in a leaf L , which is not null-homotopic in that leaf. By 'blowing' the map off of L using the transverse orientation, we can assume in addition that the map f is transverse to \mathcal{F} near the boundary of the 2-disk (so any small deformation of f is, as well, and so the singularities of the resulting pullback foliation live in the interior of the disk).

By gluing two copies of the map f together (in the case of the disk), the corresponding foliations glue together to give a singular foliation of the 2-sphere whose singularities miss the equator (since the foliation either is transverse to the boundary of the disks (in the first case), so glues together to give a foliation without singularities along the equator, or contains the boundary as a leaf (in the third case), which does the same). The Euler characteristic calculation then tells us that our foliations must have center singularities (which pair off top and bottom in the case of the disk), so we can conclude that in all of the cases of Novikov's theorem, the induced foliations on our disks and spheres must always have center singularities.

We will use this fact to show that we can always find vanishing cycles (they will in fact be the images of loops in the (possibly singular) leaves of our pullback foliations), by picking a center singularity and walking out from it. We will show that either we must eventually bump into a vanishing cycle, or we can redefine the map f to get a singular foliation with fewer center singularities. Since we can't remove them all, and we only start with finitely-many, in the second case we would eventually arrive at a contradiction.