Foliations and the Topology of 3-manifolds

Outline of class 1

Ultimately we will be focusing 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.

Books:

Tamura, *Topology of foliations*

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

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

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 $f:M^n \twoheadrightarrow N^k$ between smooth manifolds (i.e. $f_* : T_xM \rightarrow T_{f(x)}N$ 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: $f(x_1,\ldots,x_n) = (x_{n-k+1},\ldots,x_n)$. Therefore, point-inverses look, locally, like horizontal sheets. A posteriori, each leaf is an (n-k) - manifold.

So, for example, f:$\mathbf{R^3} \to \mathbf{R^1}$ 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:$\mathbf{R^3} \to \mathbf{R^1}$ 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→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:$M \setminus \Sigma(f) \to N$, which yields a foliation away from the singular set. For example, f:$\mathbf{R^3} \to \mathbf{R^1}$ given by f(x,y,z) = xyz foliates M=$\mathbf{R^3}\setminus$(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∼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→M/∼ (= (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:$\mathbf{R^2} \to \mathbf{R^1}$; one can assign a function which is constant on leaves, which has a nowhere-zero gradient (just start at -∞ and work up to 0 (without getting there) on the open ray, and start at 0 and go to ∞ on each of the clsed 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 -∞ 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:$\mathbf{R^2} \to \mathbf{R^1}$. (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:$M_{\mathcal{F}} \to M$ 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∈M, an open set set $\mathcal{U}$ and a homeomorphism $h:\mathcal{U}\to\mathbf{R^n}$ (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**

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

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 $\mathbf{R}^{\mathbf{n-k}}\times\{\mathrm{pt.}\}$of $\mathbf{R}^{\mathbf{n}}$. 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

$$t_{\alpha\beta} = (t^1_{\alpha\beta}, \ldots, t^n_{\alpha\beta}$$

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

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 coordintates 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 <u>have</u> 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 $p:\tilde{M}\to M$, 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 $h:\mathcal{U}\to\mathbf{R}^{\mathbf{n}}$. 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 $\mathcal{U}_\alpha$ each mapped homeomorphically down to $\mathcal{U}$ by p. so if we take the collection of pairs $\{(\mathcal{U}_\alpha,\, h\circ p|_{\mathcal{U}_{\mathrm{alpha}}})\}$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

$$(\mathcal{U}_\alpha,\, h\circ p|_{\mathcal{U}_\alpha}) \text{ and } (\mathcal{V}_\beta,\, k\circ p|_{\mathcal{V}_\beta})$$

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

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

(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 $\mathcal{U}_\beta$ and then back up to $\mathcal{U}_\alpha$, we can go dirctly 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) 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+m)=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 ivariant 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(mod 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 $\{k(an_0)+c \pmod 1) : 1\leq k\leq N \}S^2A$ (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_xM 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_xM. 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_xM). 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 any time 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^1_x 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^1_x 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 $\pi_1(\mathrm{M})$, since the cancatenation of two orientation-preserving loops is orientation-preserving (so G in 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 $\pi_1(\mathrm{M})$; if there are no orientation-reversing loops, then $\mathrm{G}=\pi_1(\mathrm{M})$, while if there is one γ, then any other γ' is equal to $\gamma\cdot\alpha$, where $\alpha\in\mathrm{G}$ (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 $\widetilde{\Delta}^1$ on the 2-fold (or 1-fold!) cover $\tilde{\mathrm{M}}$ corresponding to G (just take the preimage of Δ^1_{x} under the map $\mathrm{p}_* : \mathrm{T}_{\tilde{\mathrm{y}}}\tilde{\mathrm{M}}\to\mathrm{T}_{\mathrm{y}}\mathrm{M}$ for each $\tilde{\mathrm{y}}\in\tilde{\mathrm{M}}$). This distribution can then be oriented by picking a base point $\tilde{\mathrm{x}}$ and an orientation for $\widetilde{\Delta}^1_{\tilde{\mathrm{x}}}$ and then <u>assigning</u> an orientation for any other $\widetilde{\Delta}^1_{\tilde{\mathrm{y}}}$ by dragging an orientation along any path from $\tilde{\mathrm{x}}$ to $\tilde{\mathrm{y}}$. 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 $\tilde{\mathrm{y}}$, 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 <u>tell</u> if you were upstairs or down, because it is a local construction). But this is absurd, since $\mathrm{p}_*(\gamma)\in\mathrm{G}$, 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; $\pi_1(\mathrm{M})=\{1\}$ <u>has</u> no index-2 subgroup, so $\tilde{\mathrm{M}} = \mathrm{M}$!

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 $\mathrm{T_xM}$ for each $\mathrm{x}\in\mathrm{M}$) 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 $\mathrm{T_xM}$) find a k-dimensional distribution on M. Therefore, our result from last time, that a closed manifold M has a 1-dimensional distribution iff $\chi(\mathrm{M})=0$, implies <u>also</u> that a closed manifold M has a codimension-1 distribution iff $\chi(\mathrm{M})=0$. So manifolds with non-zero Euler characteristic (like the ones described last time) cannot admit a codimsion-1 foliation.

But what can we say if a manifold <u>does</u> 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 $\Delta=\mathrm{T}\mathcal{F}$ 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 <u>define</u> 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 <u>horizontal</u> 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 <u>only</u> 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 **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 distinguished 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 $X(p),Y(p)\in\Delta^1_p$ for all p are, locally, multiples of some non-zero vector field Z (so X=aZ, Y=bZ for some functions $a,b:\mathcal{U}\to\mathbf{R}$), and then a quick calculation shows

$$[X,Y](f)=aZ(bZ(f))-bZ(aZ(f))=a(Z(b)Z(f)+bZ(Z(f)))-\ b(Z(a)Z(f)+aZ(Z(f)))$$
$$=(aZ(b)-bZ(a))Z(f)$$

so $[X,Y](p)=(aZ(b)-bZ(a))Z(p)\in\Delta^1_p$ 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

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

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 it's tangent space at (1,1,0), would then be the xy-plane, which is not the same as $\Delta^2_{(1,1,0)}$ = 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 $\gamma'(t)\in\Delta^2_{\gamma(t)}$ 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 $\gamma(t)=(t,0,0)$) $\gamma'(t)=(1,0,0)\in\Delta^2_{(t,0,0)}$ 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 γ(t)=(a,t,0)) γ'(t)=(0,1,0)$\in\Delta^2_{(a,t,0)}$ 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 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 R^2. This disk inherits a codimension-1 foliation from R^2, with $\gamma=\partial\Delta^2$ as a leaf. But if we take two copies Δ^2_0, Δ^2_1 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 transverse 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 R^2, which is the union of all leaves which intersect the arc, is open in 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 enpoint 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. 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 orint 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 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 Grassmann bundle, then we can talk of <u>deforming</u> 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 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:

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

$(x^2+y^2\text{-}1)e^z=\text{-}c^2<0$ implies $x^2+y^2=1\text{-}c^2e^{-z}<1$, so this leaf lives inside the cylinder; as $z\to\infty$, $x^2+y^2\to1$, and as $z\to\text{-}\infty$ the right hand side becomes negative, so the leaf closes up by hitting the z-axis and does not continue on down,

$(x^2+y^2\text{-}1)e^z=c^2>0$ implies $x^2+y^2=1+c^2e^{-z}>1$, so these leaves live outside the cylinder; as $z\to\infty$, $x^2+y^2\to1$, and as $z\to\text{-}\infty$, $x^2+y^2\to\infty$, 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 diresction, 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 $D_2\times R^1$ 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 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 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 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 in one 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 $R^3 \cup \{\infty\}$), 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 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 transverse 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 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 $L\cup\{(\text{z-axis})\cup\infty\}$ 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 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 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^3open 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 $D_1,\ldots,D_n$ of Dehn twists along non-separating curves such that $D_n \circ \ldots \circ D_1(\gamma')$ 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 $f_1(\gamma)=D_n \circ \ldots \circ D_1 \circ f(\gamma)$ 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, gamprsp in F_g meeting one another transversely (in a finite number of points), then there is a sequence of simple closed curves $\gamma=\gamma_1,\ldots,\gamma_n=\gamma'$, 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 it 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, of 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 $\mathrm{F_g}\backslash\gamma,\mathrm{F_g}\backslash\gamma'$). 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, gamprsp 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 $\mathrm{k{\circ}h^{-1}}$ 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 $\mathrm{D_k}$ 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.

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 <u>same</u> 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 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 <u>all</u> closed 3-manifolds M. And about the only thing that fits that bill is that $\chi(M)=0$ (which a foliation <u>does</u> 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) an 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, $\widetilde{M}$, is homeomorphic to 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 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 <u>all</u> 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. π_1(M) is infinite,

2.π_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 $\pi_1(\mathrm{L})\rightarrow\pi_1(\mathrm{M})$ 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 $\widetilde{\mathrm{M}}$, then the commutativity of the diagram below easily shows that any leaf $\widetilde{\mathrm{L}}$of $\tilde{\mathcal{F}}$ π_1-injects into $\widetilde{\mathrm{M}}$, i.e., it is simply-connected. But the only simply-connected surfaces are S^2 and 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 <u>other</u> cases that every leaf of $\tilde{\mathcal{F}}$ is a plane. Then the following result becomes very interesting.

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

Setting N=$\widetilde{\mathrm{M}}$ 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 $\widetilde{\mathrm{M}}\cong$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. π_1(M) is finite,

2. π_2(M)$\neq$0, or

3. for some leaf L of $\mathcal{F}$, the inclusion-induced homomorphism is <u>not</u> 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.

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 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 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 do 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 alot 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→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 its 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.

We can apply this reasoning, for example, to loops γ in leaves (which are formally maps f:[0,1]→L with f(0)=f(1)). Over (the image of) this loops 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).

On 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→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 <u>not</u> 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 - its 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 if a singular (although by general nonsence 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 <u>embedded</u> 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 (π_2(M)$\neq$0), <u>either</u> $\mathcal{F}$ has a vanishing cycle <u>or</u> 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$\cong\mathrm{S}^2\times S^1$, every leaf of $\mathcal{F}$ is a 2-sphere, and the foliation consists (up to homeomorphism) of leaves $\mathrm{S}^2\times\{\text{pt.}\}$.

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 $\mathrm{S}^2\times S^1$; 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

($\mathbf{RP^3}\backslash$(3-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, $\mathrm{M}\cong\mathrm{S}^2\times S^1$.

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 $\mathrm{S}=\mathrm{D}^+\cup\mathrm{D}^-$, 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, $\mathrm{D}^+\cap\mathrm{D}^-=\gamma$ 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 $\mathrm{x}\notin\mathcal{U}$ 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 $\mathrm{x_n}$ in $\mathcal{U}$ limiting down on x. In particular, in some distinguished chart about x, the plaques containing the $\mathrm{x_n}$ 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 $\mathrm{x_n}$ 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 $\mathrm{x_n}$) 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 $\mathrm{x_n}$ and x in an open set, so there is a first point z along this arc going from $\mathrm{x_n}$ 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 $\mathrm{f}:\mathrm{S}^2\times[0,\epsilon]\to\mathrm{M}$ which maps 2-sphere levels injectively into leaves of $\mathcal{F}$, for which $\mathrm{z}\in\mathrm{f}(\mathrm{S}^2\times\{0\})$. 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.

Bu 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 $S^2_r=S^2_s$ 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 pict any point in S^2_t 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 $S^2_r=S^2_s$, 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 to 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^2_ϵ are defined for all $t \geq 0$, and have finite length. But then it is easy to see that the map $f{:}S^2_\epsilon \times [0,\epsilon]$ given by $f(x,t)=$ the point along the trajectory from x in the leaf S^2_t 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 $S^2_\epsilon \times 0$ 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 <u>only</u> 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 <u>has</u> 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 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 t_0, ϵ.

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, since we can flow both backwards and forward from the resulting disk D_{t0}, 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.

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 $x \in \mathrm{int}(D_\epsilon)$ 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 $t_n \to 0$.

But this plaque sits on the leaves L_{tn}, and in fact we must eventually have $y \in D_{tn}$, 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 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 $D_\epsilon$in the same amount of 'time'. Therefore the trajectory through x (since it passes arbitrarily close to y infinitely-often) hits $D_\epsilon$infinitely-often. Therefore there is a first time t_0 in which it happens.

We therefore have for this time t_0 that $D_{t0} \cap D_\epsilon \neq \emptyset$, 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_{t0}, then it also contains γ_{t_0}, implying the $D_\epsilon \cup A$ is not embedded; so we have that $D_\epsilon \subseteq D_{t0}$.

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 $f(\partial D_\epsilon \times [t_0,\epsilon])$ (see figure below).

But if we look at the other annulus $B = D_{t0} \setminus \mathrm{int}(D_\epsilon)$, 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 $t\to 0$, 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, 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 entir 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 $T_0 \times \{pt.\}$, 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 $T^2 \times \{pt.\}$, and turbulize along a loop $\{pt. \}$ 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 <u>begin</u>: 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 codmension-1 foliations are the natural setting for something like a Morse theory of functions $f{:}F^2 \to M^3$ to make sense.

A **Morse function** is a smooth map $f{:}F^2 \to R^3$ from a surface F to Euclidean space so that the critical points of $p \circ f : F \to R^1$(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 $f_*(T_xF)+T_{f(x)}\mathcal{F}=T_{f(x)}R^3$).The point is that this fails ($f_*(T_xF)=T_{f(x)}\mathcal{F}$) 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 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 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; one 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 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}(horizontal plane)), at least, off of the critical points. These foliations one 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 <u>other</u> 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 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 <u>vector</u> <u>field</u> (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 <u>know</u> 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, $\pi_1(\mathrm{M})$ 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 $\mathrm{f{:}D_2}\to\mathrm{M}$ such that $\mathrm{f}|_{\partial\mathrm{D}}=\gamma$ 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 $\mathrm{S}^2\times S^1$, 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 R^2, and so) has π_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:$D_2 \to$M such that $f|_{\partial D}=\gamma$ is a loop in a leaf L, which is <u>not</u> 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 <u>fewer</u> 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.

Foliations and the Topology of 3-manifolds

Outline of class 11

Last time we saw how to use a map f from a surface F into our 3-manifold M to pullback our foliation $\mathcal{F}$ on M to give a singular foliation $f^*(\mathcal{F})$ on F, and we noticed that in the cases we will consider, F=D_2,S^2, the singular foliation must always have a center singularity (as well as possibly some saddle singularity). If we start walking away from such a center, and consider the (possibly singular) leaves of $f^*(\mathcal{F})$ that we pass through, initially they are loops (since this is true in the local picture around a center critical point that we drew last time) that are null-homotopic in their leaves. What we must consider is what happens the <u>first</u> time we encounter something <u>other</u> than such loops. If we think about it for a bit, it's not hard to see that one of five things must happen:

(1): We encounter a loop γ which is <u>not</u> null-homotopic in its leaf. Since on one side of this loop (the one we approached it from along f(F)) the normal fence hits nearby leaves in (loops closely approximated by the leaves of our singular foliation, hence) loops null-homotopic in their leaves, this loop γ is therefore a vanishing cycle, and we can stop, having found what this construction was trying to find for us.

(2): We run into a non-compact (possibly singular) leaf of $f^*(\mathcal{F})$. But this is actually impossible. Such a loop, since F is compact, must therefore eventually pass arbitrarily close to itself, and so can be short-circuited to a loop γ' transverse to our singular foliation, as we have done in the past. But since the loops around the center (which, you should notice, are null-homologous) are limiting on this non-compact leaf, they pass arbitrarily close to it, and so intersect the transverse arc we used to short-circuit our leaf. But since the singular foliation can be transversely oriented (and hence oriented, by uniformly turning the transverse orientation to one side or the other), the loops γ and γ', when we give them orientation, always intersect one another with the same sign (see figure below); consequently, their homological intersection number is non-zero, implying both are not null-homologous, which in a disk or sphere is of course absurd. So this case can't happen.

(3): We run into another center singularity. We'll deal with this case after the next two; in this case we shall see we have a 2-sphere, and have found our one exception to the rule of always finding vanishing cycles.

(4) We run into the boundary of the disk. But then we must be in the case that the boundary is a leaf of our singular foliation (if the singular foliation were transverse to the boundary, then the loops surrounding our center couldn't be limiting on the boundary (exercise - it's because they don't hit the boundary themselves). But in that case by assumption this leaf of the singular foliation is not null-homotopic in its leaf, so we are really in case (1), and we are done.

(5): We run into a saddle singularity. Since the part of the singular leaf containing the saddle that we have in fact encountered must be compact (otherwise we're in case (2)), we must have in fact one of the following two pictures:

We can't have our center circles running into 3 or more of the corners of the singular leaf (because an arc joining two of the arcs emanating from the saddle forms a (separating) loop, which must therefore be joining an adjacent pair of such arcs), so there are only these two pictures (we can't run into two saddles at once, since they map into different leaves of $\mathcal{F}$, so can't be in the same singular leaf of $f^*(\mathcal{F})$). Leaving the first picture alone for a minute, let's see how we can deal with the second case.

We know that all of the circles inside of the disk that the singular leaf cuts off are null-homotopic in their leafs of $\mathcal{F}$. If the loop γ in the singular leaf is not null-homotopic in its leaf, we are done: it is a vanishing cycle by the argument in case (1) above. If it is null homotopic in its leaf, we then cancel the center and saddle singularity as in the sequence of pictures below:

We first redefine f on the disk (containing the center) cut off by γ, so that it instead maps into the leaf containing γ (using the fact that it is null-homotopic). This map is no longer Morse; all of the points in the disk are degenerate critical points. But if we 'smooth' off the map in a neighborhood of this disk (basically by starting at the saddle point and pushing it own into nearby leaves (going the direction that nearby non-singular leaves lie) in 'concentric arcs' (see figure)), then we get a new Morse function, identical to the first away from a neighborhood of the disk, which now has one fewer center and saddle singularities.

In the first case, we don't deal with this center singularity right away. Instead we go find a different one to deal with. There can be only finitely-many of these 'inside-out figure-8's' (since there are only finitely many saddles), and so if we look at the disks that each half of the figure-8's bound (they bound disks since they are embedded and are in a disk (or sphere)), there is an 'innermost' one whose interior doesn't intersect any of the figure-8's (each such disk contains fewer figure-8's than our original one (it doesn't contain the one it bounds!); continue by induction). On this disk we find a singular foliation which (if we ignore the 'corner' it has from running into its saddle on the boundary) actually satisfies one of the conditions we used in our original Euler characteristic count; its boundary is a leaf. So there must be a center we can carry out the singularity in this disk, and we can carry out the analysis we have given above, to either find a vanishing cycle, or (possibly by cancelling with the saddle that is in the boundary of our small disk) cancel the center with a saddle, reducing the number of both. We can't run into the figure-8 problem, because they are all outside of our small disk - we would run into the boundary first, which is a second-case example of a saddle.

So in every case (except (3)) we can always either find a vanishing cycle or cancel a center with a saddle. Since we have never tampered with the singular foliation at the boundary (in the disk case), we are always assured of having a center singularity, so if we don't find a vanishing cycle, it must be the case that we run out of saddles; we can deform our singular foliation so that it only has center singularities (in case (3), we leave it as an exercise to verify that the singular foliation actually already <u>has</u> no saddles - the foliation is a disk foliated by concentric circles, with the boundary crushed to a point. Therefore, our surface, in particular, is a 2-sphere!). If we <u>still</u> don't run into a vanishing cycle as we move out from a center, our only recourse then is case (3); so we are dealing with a 2-sphere such that $f^*(\mathcal{F})$ is a singular foliation by latitudes (see figure). What we will show is that <u>if</u> we assume that $\pi_2(L)=0$ for all leaves of $\mathcal{F}$, then this map $f:S^2\rightarrow M$ can be extended to a map of the 3-ball inot M. Since f by assumption (in the case of a 2-sphere) is non-trivial in $\pi_2(M)$, this is a contradiction. Therefore it must be the case that $\pi_2(L)\neq 0$ for some leaf of $\mathcal{F}$. This leaf is then (either S^2 or $\mathbf{RP^2}$ (these are the only surfaces with non-trivial π_2), so by orientation assumptions is) S^2. So $\mathcal{F}$ has an S^2 leaf, so by Reeb Stability, $M=S^2\times S^1$, and every leaf is a 2-sphere; this was our one exceptional case in Novikov's theorem.

We will build the extension of f to the 3-ball in steps. The main fact we will use is that the homotopy disks that the latitude loops bound (which exist by the assumption that we didn't run into a vanishing cycle before we ran into the second center) can be lifted in a continuous fashion to homotopy disks bounding nearby loops of the singular foliation. In particular, if we imagine our circles parametrized by [-1,1], with -1 and 1 representing the center singularities, then for every $t \in (-1,1)$ there is an $\epsilon=\epsilon(t)>0$ so that the singular disk bounded by the loop γ_2 lifts to leaves containing the loops within ϵ of t (by the standard picture of a center, the loops close to the center lift to null-homotopic loops all the way to the leaf containing the center). We can also assume that $\epsilon(t)$ is small enough so that the loops we lift in the normal direction are freely homotopic to the corresponding loops of the pullback foliation $f^*(\mathcal{F})$, so we can in fact think of these lifts as lifting to the nearby γ_2's. Since the interval [-1,1] is compact, it can therefore be covered by a finite number of these intervals, so we can assume we have chosen intervals $[-1,s_1),(r_1,s_2),\ldots,(r_{n-1},s_n),(r_n,1]$ so that $r_k<t_k<s_{k+1}$ for all k, and the loop at level t_k lifts as above to both r_k and s_{k+1}. Now pick x_k such that $s_k<x_k<r_k$. Then the disk at level t_k lifts down and up to $|rmx_{k_1}$ and x_k (see figure below). The disks down below and above the loop at x_k lift up and down to two disks whose boundaries are γ_{xk}, so gluing the disks together, we get a map of a 2-sphere into the leaf containing γ_{xk} (see figure). Because $\pi_2(L)=0$ for every leaf, the map of this sphere extends to a 3-ball, for each k. But then gluing all of these maps of the 3-balls together with the lifts of the 2-disks from x_k to x_{k+1} for each k, we get a map of a 3-ball (a bunch of 3-balls glued end to end form a 3-ball) which on the boundary is our original map f (see figure). So f is null-homotopic, a contradiction (as desired!).

So we have now finished the <u>first</u> part of our outline of the proof of Novikov's theorem. We have found (except in the exceptional case $M=S^2\times S^1$) a (singular) vanishing cycle γ under each of the three hypotheses of Novikov's theorem. By general nonsense (transversality theory), since we are dealing with a loop in a surface, we can homotope γ so that it is immersed and transverse to itself. It is still a vanishing cycle: it is still a non-limit cycle (on the same side) and non-trivial in its leaf, since being non-limit (and non-null-homotopic) is invariant under free homotopy, and the lifted loops are still null-homotopic in their leaves, since this is again unchanged by free homotopy.

What we will do now is show how to exchange this immersed vanishing cycle for an embedded one; this will finish our proof of Novikov's theorem, since we have already shown how an embedded vanishing cycle sits in the torus leaf boundary of a Reeb component. We will find this new, embedded, vanishing cycle in two steps. First we will show that we can find an immersed vanishing cycle γ (in a possibly different leaf of $\mathcal{F}$) whose lifts to nearby leaves along its normal fence are all <u>embedded</u>; then we will show how to use this new vanishing cycle and its lifts to find an embedded vanishing cycle.

The first part will utilize two facts. The first is that if γ is a loop in a leaf with n self-intersections, then all of the lifts of γ close enough to γ have at most n self-intersections. This is because the only way for a lift to intersect itself (over a point where γ doesn't) is for it to 'catch up' to itself (see figure), so one of the lift must bump into γ again first - but we can clearly avoid that for a short time.

The lifts of γ to the nearby leaves are null-homotopic loops which are immersed transverse to itself. What we are going to try to do is to 'unwrap' these lifts to make them embedded. To do that we need to understand what a null-homotopic loop in a surface really looks like. With that in mind, we have the following

Proposition: If $\gamma \subseteq$F is a null-homotopic loops in a surface F, immersed transverse to itself, then γ can be transformed to an embedded (null-homotopic) loop (which therefore bounds a disk) by the following three moves:

The reader familiar with knot theory will recognize that these moves are the 'immersed' versions of Reidemeister's moves, by which a knot can be transformed to any other knot isotopic to it.

Next time we will show first how this proposition allows us to find a vanishing cycle with embedded lifts, and how to then find an embedded vanishing cycle. Then we will show how to prove the proposition; it should come as no surprise that the proof will be by induction (on the genus of F and the number of self-intersections of γ).

Foliations and the Topology of 3-manifolds

Outline of class 12

OK, so that wasn't the greatest lecture in the world. Or even close. Here's what I should have been saying.

We are trying to use an immersed vanishing cycle to find an embedded one. Since the lifts of an embedded vanishing cycle are embedded (null-homotopic) loops, one way to get halfway there is to find a vanishing cycle whose lifts are embedded. This is what we do first.

Suppose γ=γ_0 is our immersed vanishing cycle, with n self-intersections, and consider the family of lifts γ_2, 0≤t≤ϵ, of γ to nearby leaves (where ϵ is small enough so that all of the lifts have at most n self-intersections). If all of the lifts sufficiently close to γ_0 are embedded, then we are done. Otherwise, some nearby leaf (γ_ϵ, say) has some self-intersections. By the proposition, we can deform γ_ϵ to an embedded loop by a sequence of 'Reidemiester' moves. Since the third one does not change the number of self-intersections of γ_ϵ, after a (possibly empty) sequence of type 3 moves, we must have a type 1 or 2 move. This move then gives us an embedded disk (whose boundary δ_ϵ consists of 1 or 2 subarcs of γ_ϵ) which, if we let it (i.e., its boundary) flow back under the type 3 moves, gives an immersed disk D_ϵ bounded by an (immersed) loop made up of 1 or 2 (immersed) subarcs of γ_ϵ (see figure below). Since the subarcs live in the normal fence over γ, they each can be pushed down to subarcs of each of the γ_2's. As in the proof of Reeb Stability, the immersed disk D_ϵ can be flowed down in the direction of γ to disks D_t bounded by loops δ_t, made up of the arcs we found above, at least for t's near ϵ. The collection of points (in the normal fence over (one of the possibly two) intersection points for which these arcs give loops which are null-homotopic in their leaves is open (since the null-homotopies flow up and down). If this set is not all of the points in the fence, then there is a first point not in the set. This point is still in a loop (if the arcs don't close up, then the lifts of the resulting path to nearby leaves are also not closed), and this loop is therefore a vanishing cycle. This vanishing cycle (once we smooth the corners coming from the 1 or two intersection points of $\gamma_\epsilon$in δ_ϵ - see figure) has fewer self-intersections (at least 1 or 2 fewer) than γ_0. So we would be done by induction on the number of self-intersections of the vanishing cycle (since we can't continue reducing the number of self-intersections).

But if the set of is all of the set, then the arcs push down to a null-homotopic loop in the leaf containing γ_0, then (see figure) we can use the null-homotopy to deform $\gamma_0$to a loop with fewer self-intersections. In the type one case, we merely erase the loop (and smooth out the resulting corner; the disk can be used to deform the loop to the corner point. This loop has at least one fewer self-intersection than the original (the corner we smoothed out).

In the type 2 case, we can use the null-homotopy to interchange the two arcs; then pushing them a little bit further apart, we can reduce the number of self-intersections by two. This fails in exactly one case: when the two endpoints of the 'digon' are actually identified (see figure). Then the interchange actually gives the same loop back. But in this case The loop upstairs (as well as the one downstairs) is the square of a single loop (consisting of one of the two arcs - its endpoints are identified). Upstairs, since the square is null-homotopic, and since the fundamental group of the leaf is torsion-free (this is true of any orientable surface), the loop itself is null-homotopic. Downstairs, since the square is not null-homotopic, the loop is also not null-homotopic. So either the null homotopy pushes all the way down, so the loop downstairs is a vanishing cycle (with at least one fewer self-intersection), or it doesn't; at the point it stops, we once again get a vanishing cycle, with fewer self-intersections.

So in all cases, if a lift has self-intersections, we can find a new vanishing cycle with fewer self-intersections. Therefore by induction we can find a vanishing cycle whose lifts are embedded.

It is worth noting that we can in fact have a non-embedded vanishing cycle whose lifts are all embedded; we can build one by taking the meridian loop in the boundary of a

Reeb component and drag a piece of it around the torus in the longitudinal direction until it intersects itself; see the figure.

We will now use this new 'improved' vanishing cycle to find an embedded vanishing cycle. The idea is that since we now have embedded, null-homotopic, lifts, they bound embedded disks D_t in their leaves. In particular this is true of γ_ϵ; it bounds D_ϵ. What we will do is look at the intersection of D_ϵ with the normal fence A over γ. $A \cap D_\epsilon \subseteq D_\epsilon$ is a closed subset of D_ϵ which consists of a finite number of arcs of the loops γ_2, $t<\epsilon$(because the fence is transverse to the leaves of $\mathcal{F}$, and D_ϵ is contained in one; there are finitely-many, since the disk is compact, so can hit a normal fence at only a finite number of levels). There can be no closed loops of intersection; if there were, D_ϵ flows down along the fence to a disk contained in itself, bounded by γ_{t_1}, $t_1<\epsilon$; but letting that subdisk flow down we get another subdisk (contained in it) bounded by γ_{t_2}, $t_2<t_1$; continuing inductively, we could then in fact conclude the D_ϵ hits (in fact contains!) an infinite number of the γ_t's, a contradiction. We therefore have a picture like in the figure below.

These arcs cut the disk up into pieces (subdisks, since no piece can have disconnected boundary (because there are no loops)). Each piece therefore gives us a loop in the disk. There are now two things to verify, to finish our proof:

(1): If all of these loops push down all the way to γ_0 as closed loops, then the image of one of them is not null-homotopic. This implies, in the standard way, that one of them pushes down to a vanishing cycle at some level (if it doesn't push as a loop all the way down, then at some point along the way it becomes a vanishing cycle).

(2): If we take the loops consisting of the boundaries of a small neighborhood of the union of these arcs (and the boundary - see figure above), then all of them are embedded loops (for as long as they are loops). Therefore, whichever one, when pushed to a vanishing cycle, is in fact an embedded vanishing cycle, completing our proof.

The first one is not too hard; it follows from the fact that the boundary of the disk (i.e., γ_ϵ) can be written as a composition of these smaller loops (with change-of-basepoint arcs attached, so they are all based at the same point). This follows quickly by induction on the number of arcs in the disk; see the figures below. Each time we add one, we build the composition by induction, by adding one piece of the arc at a time (so we are always basically doing the same thing - subdividing one loop into two).

The second one is a little tricky. As a matter of fact, it is SO tricky that I don't know how to prove it. I still think it's true, but the proof I had in mind goes sour at one point (I don't know that what ought to happen does). BUT as it happens, I know of a different way to show how to pass from a vanishing cycle with embedded lifts to an embedded vanishing cycle - it's amazing what deadline pressures will occasionally do to increase brain function. The idea is to show first that the leaf containing the vanishing cycle is a torus; then by using a little knowledge of what loops in a torus look like, we can pretty easily trade our vanishing cycle in for an embedded one. We'll do this next time.

Foliations and the Topology of 3-manifolds

Outline of class 13

Today we will finish our proof of Novikov's Theorem (modulo the result about deforming null-homotopic loops in surfaces - we won't do this in class, but I will write it up (soon) for these notes). We have shown that we can find (in all of the cases of the theorem) an immersed vanishing cycle γ in a leaf L, whose lifts to nearby leaves are embedded null-homotopic loops γ_t (and therefore bound disks D_t in the leaves containing them). We will show now how to replace γ with an embedded vanishing cycle, which will complete our proof, since we have shown that embedded ones are contained in torus leaves of Reeb components.

The basic outline is to show, first, that the leaf L containing our vanishing cycle is a torus. Then it is a relatively simple matter to show, using some facts about what loops in a torus look like, to trade our vanishing cycle with an embedded one (which the original is a power of).

The first part can be seen by mimicking (to the extent possible) our construction of a Reeb component in the case that the vanishing cycle is embedded. The point is that most of the construction relied on the fact that the null-homotopic lifts were embedded, and not that the vanishing cycle was. We start by building, as in that construction (from now on called 'the embedded case'), a map $f:D_\epsilon\times(0,\epsilon]\rightarrow M$ by starting at one of the lifts D_ϵ and flowing along the trajectories of the transverse foliation from one disk to nearby ones. For the same reasons as the embedded case, this map is defined for all $t>0$; if not, there is a first t that the map doesn't extend for all of the disk; but since there is a disk at that level, it can either be flowed up to give us a way to extend past t (a contradiction), or the lifted disk does not agree with the disk already there (see figure). But then the two disks we now have glue together to give a sphere leaf of $\mathcal{F}$, so the foliation consists of spheres, again a contradiction (there's no leaf to contain a vanishing cycle; all loops are null-homotopic). This map is a local homeomorphism; this is because the disks D_t are embedded, so (since we are flowing along the transverse vector field, which can't send two nearby points to the same point (by the uniqueness of solutions of differential equations). So small (cubical) neighborhoods of points in $D_\epsilon\times(0,\epsilon]$ are mapped injectively by f.

But again, this map cannot be extended to $t=0$; if it could, it would give a map of a disk into the leaf L, restricting to the loop γ on the boundary, implying it is null-homotopic, a contradiction. Then, again as in the embedded case, some trajectory out of a point $x\in D_\epsilon$

becomes infinitely long as t tends to 0, and therefore (again, as in the embedded case!) this trajectory passes through $D_\epsilon$infinitely-often as t tends to 0.

This is the point where we first break away from our embedded case; in that case, the first time t_0 that the trajectory returned, we had $D_\epsilon\subseteq\mathrm{int}(D_{t0})$; this was because D_ϵ did not intersect any of the loops γ_t for $t<\epsilon$. Here we cannot conclude this, because, since the loop γ is in general not embedded (but the lifts are), the loops gamteesp <u>must</u> be intersecting one another. However, because D_ϵ is compact, it can intersect only finitely-many of the γ_t's. If it intersected infinitely-many, then (since the normal fence over γ can be covered by finitely-many distinguished charts) D_ϵ would intersect a chart in infinitely-many plaques, giving an infinite set in D_ϵ with no limit points, a contradiction. Therefore, there is a (in fact, infinitely-many) disks D_{tn} with $D_{tn} \cap D_\epsilon \neq \emptyset$, but $D_\epsilon\cap\partial D_{tn}=\emptyset$. Consequently, $D_\epsilon\subseteq\mathrm{int}(D_{tn})$. In fact, we can (inductively) arrange that $D_{tn}\subseteq\mathrm{int}(D_{t_{n+1}})$ for a sequence of t_n's tending to zero; just let (inductively) D_{tn} play the role of D_ϵ in the argument above.

Now we rejoin our embedded proof. Look at our function f restricted between these two disks; $f{:}D_\epsilon\times[t_0,\epsilon]\rightarrow M$. By the above, $f(D_\epsilon\times\{\epsilon\})=D_\epsilon\subseteq f(D_\epsilon\times\{t_0\})=D_{t0}$, so if we create a new space $X=D_\epsilon\times[t_0,\epsilon]/$, where (x,t_0) (y,ϵ) if $f(x,t_0)=y=f(y,\epsilon)$, then f defines a map from X to M (since we have identified point which get identified under f). X is, as in our embedded proof, a solid torus whose boundary $\partial x=A_0\cup A_1$ (see figure), a union of two annuli, one $A_0=\partial D_\epsilon\times[t_0,\epsilon]$ mapping into the normal fence over γ, and the other A_1 mapping (injectively!) to $D_{t0}\backslash D_\epsilon$.

For large enough n, the points of $D_{tn}\backslash D_\epsilon=f(A_1)$, as in the embedded case, come which have limiting values as t tends to 0, although the argument is more difficult here than in that case (since there we could argue that X was <u>embedded</u> in M (under the map f)). We do know, as in the embedded case, that if a trajectory through a point of $f(A_1)$ becomes infinitely long as t tend to 0 (which is the only way for a point not to have a limiting value), then it must pass through $D_\epsilon$infinitely-often. In particular, it must pass through D_ϵ <u>after</u> time t_n, i.e., after passing through $f(A_1)$. We will show, however, that this can't happen.

Because the trajectory hits D_ϵ <u>after</u> it hits D_{tn}, we can short circuit the trajectory when it returns to D_ϵ, and turn it into a loop δ transverse to the foliation and intersecting A_1, hence the (singular) torus $f(\partial X)$; see figure. This loop misses (the image of) A_0 (since trajectories hitting it have limiting values (just apply the argument for large enough n so that the normal fence from γ to γ_{t_n} does not hit D_ϵ- the usual argument shows that $D_\epsilon$eventually doesn't intersect the γ_t's, for small t)), and so all of the intersection points of δ with $T=f(\partial X)$ occur along A_1as intersections of the trajectory. But these intersections always occur with the same sign (in the sense of intersection numbers on homology), so the intersection number of δ with T is non-zero, implying that both are non-trivial in homology. But this is absurd; T bounds the (singular) solid torus f(X) (another way to say this is that the homology class of T is the image under f_* of the homology class of ∂X, which is zero), so is trivial in $H_2(M)$.

But now the annulus A_1 can be mapped, by the limiting values of the trajectories, into the leaf L containing γ; furthermore, since $f(A_1)$ is in a leaf of $\mathcal{F}$ and so transverse to the trajectories, the same proof as outlined above show that this limiting map is a local homeomorphism. Further, since ∂A_1 consists of two lifts of γ, they both map to γ and so we can glue the ends of A_1 together to give a torus, which (the induced function) f maps into L as a local homeomorphism. We therefore get an immersion of a torus T into L, but since T and L have the same dimension, f is therefore a covering projection. Since its image is compact (hence closed); if it isn't all of L, then the image has a boundary point. This point comes from a point upstairs, but then a small neighborhood of this point maps to a neighborhood of the point downstairs, contradicting the fact that it is a boundary point of the image. Therefore, Every inverse image of a point downstairs has a neighborhood mapping injectively; but since T is compact, a point can have only finitely-many preimages (otherwise the inverse image has a limit point, and the map couldn't possible be locally injective at that point). From here it is easy to show that the map has the local triviality property of a covering map; we leave it as an exercise.

So we now have a covering map form a torus T our leaf L containing the vanishing cycle; but the only (oriented - M is orientable and $\mathcal{F}$ is transversely-oriented) surface covered by a torus is a torus; just argue by Euler characteristics ($\chi(L)=n\chi(T)=0$ for some n). So the leaf containing γ is a torus.

Now that we know our leaf is a torus, it is a pretty simple matter to find an embedded vanishing cycle. To do it, we need to look at (non-null-homotopic) loops in a torus T. If we think of T as the quotient of its universal covering R^2 by the covering translations given by $h(x,y)=(x+n,y+m)$ for $n,m\in\mathbf{Z}$. If we take an essential loop γ in T, and assume it is based at $x_0=$ the image of (0,0)), then when we lift it up to R^2 based at (0,0) we get a path $\tilde{\gamma}$ whose other endpoint (also maps to x_0 so is) (a,b) for some $a,b\in\mathbf{Z}$. If we instead take the straight line segment ℓ from (0,0) to (a,b), then they are homotopic rel endpoints (since together they form a (null-homotopic) loop), and this homotopy projects to T to give a (based) homotopy from γ to the image of ℓ (call it τ).

But if we write d=gcd(a,b), then ℓ passes through (a/d,b/d), and the line segment from (0,0) to (a/d,b/d) projects to an embedded loop in T (otherwise two points of the lift differ by integer coordinates, so (by translating the first point to (0,0)), we get another integer point on the line segment between (0,0) and (a/d,b/d) (contradicting (exercise!) the fact that d is the gcd). Further, $\gamma=\tau^d$ in $\pi_1(T)$, since τ^d lifts in R^2 to the line segment from (0,0) to (a,b).

But the loop τ is a vanishing cycle! We see this in two steps.

First, τ is a non-limit cycle. This is because γ is a non-limit cycle, so τ^{d} is, because it is homotopic to γ. But if τ were not a non-limit cycle, then in the normal fence over τ we see lifts of τ which don't close up, arbitrarily close to τ (see figure). But for any point in this fence which are not fixed by the return map defined by these lifts, it is also not fixed by the d^{th} iterate of the return map (which is actually the return map of τ^{d}); this is just a fact about order-preserving homeomorphisms of an interval. If f(a)<a, then, inductively, $\mathrm{f}^{\mathrm{n}+1}(\mathrm{a}) < \mathrm{f}^{\mathrm{n}}(\mathrm{a})$, since otherwise (see figure) the first time the sequence reverses itself ($f^{n+1}(a) = f^n(a)$ or $f^{n+1}(a) > f^n(a)$) we can easily find two points mapped to one under f.

But since τ^{d} is not null-homotopic, τ is, and since the lifts of τ^{d} are loops null-homotopic in their leaves (since they are homotopic (in the leaves)to the lifts of γ - just lift the homotopy downstairs), the lifts of τ are null-homotopic. This is because the lifts of the d^{th} iterate of τ are the same as the d^{th} iterates of the lifts of τ, and the fundamental groups of the leaves are torsion-free (as is the fundamental group of any orientable surface). So since the d^{th} iterate of a lift is null-homotopic, the lift itself is null-homotopic.

Therefore the loop $\tau \subseteq \mathrm{T}$ is an embedded vanishing cycle, and we are done!

Our next task will be to improve on Novikov's result (2); if M admits a foliation without Reeb components and $\mathrm{M} \neq \mathrm{S}^2 \times S^1$, then $\pi_2(\mathrm{M})=0$. The improvement is Rosenberg's theorem (proven not long after Novikov's proof appeared) - under the same conditions, M is irreducible - any 2-sphere embedded in M bounds and embedded 3-ball. The proof parallels, in many ways, the construction of immersed vanishing cycles - we simply take care to insure we maintain an embedded sphere as we follow through it.

Foliations and the Topology of 3-manifolds

Outline of class 14

Today we prove:

Theorem (Rosenberg): If M is a closed orientable 3-manifold, not $\mathrm{S}^2 \times S^1$, and M admits a a transversely-orientable foliation $\mathcal{F}$ without Reeb components, then M is —bf irreducible: every (smoothly) embedded 2-sphere in M bounds a 3-ball.

We shall prove it by following our proof of Novikov's $\pi_2(\mathrm{M}) \neq 0$ implies $\mathcal{F}$ has a vanishing cycle argument. The only added wrinkle in this setting is that we must always insure that when we move our 2-sphere S (as we do in Novikov's argument) we always still have an embedded 2-sphere. We will actually do this last part by surgery, cutting our 2-spheres into (simpler) 2-spheres; then we prove that if after surgery the two 2-spheres bound balls, then the original 2-sphere bounds a ball.

We begin however with a simplification. Under the hypotheses of the theorem, Novikov's theorem implies that all of the leaves of $\mathcal{F}$ are π_1-injective in M, and since none of them are 2-spheres, they each therefore have universal cover R^2. Therefore if we lift $\mathcal{F}$ to a foliation $\tilde{\mathcal{F}}$ of the universal cover $\tilde{\mathrm{M}}$ of M, then for every leaf $\tilde{\mathrm{L}}$ of $\tilde{\mathcal{F}}$ (which projects as a covering map to a leaf L of $\mathcal{F}$), the composition (using the obvious maps) $\pi_1(\tilde{\mathrm{L}}) \to \pi_1(\tilde{\mathrm{M}}) \to \pi_1(\mathrm{M})$ is (equal to $\pi_1(\mathrm{ltil}) \to \pi_1(\mathrm{L}) \to \pi_1(\mathrm{M})$, which is) injective, so $\pi_1(\tilde{\mathrm{L}}) \to \pi_1(\tilde{\mathrm{M}})=0$ is injective. So every leaf of $\tilde{\mathcal{F}}$ is simply-connected (and not equal to a sphere), hence is a plane. So $\tilde{\mathcal{F}}$ is a foliation of $\tilde{\mathrm{M}}$ by planes.

What we will actually show is that a simply-connected (possibly (certainly!) non-compact) 3-manifold which is foliated by planes is irreducible. This suffices to prove our theorem, since:

Proposition: If the universal cover of a 3-manifold M is irreducible, them M is irreducible.

Proof: This is completely analogous to our proof that a null-homotopic, embedded loop in a surface bounds an embedded 2-disk. We lift our 2-sphere to a collection of (disjoint) 2-spheres in $\tilde{\mathrm{M}}$; there they each bound balls. Then we show that the balls they bound are all disjoint from one another, therefore any one of them maps down injectively to M to give a 3-ball bounded by our 2-sphere. We leave the details as an exercise.

So we now prove our theorem under the assumption that M is simply-connected and $\mathcal{F}$ is a foliation by planes. Given an embedded 2-sphere S in M, we can wiggle it slightly to make the inclusion map i:S$\to$M a Morse function w.r.t. the foliation $\mathcal{F}$ (since S is compact, hence can be covered by finitely-many distinguished charts). This gives us, as in Novikov's theorem, a singular foliation $\mathrm{i}^*\mathcal{F}$ on S. We proceed now, as before, to try to cancel all of the saddle singularities against the centers (there are more centers than saddles, by Euler characteristic considerations). But before starting this, we need:

Lemma: All (singular) leaves of $\mathrm{i}^*\mathcal{F}$ are compact, in the leaf topology.

Proof: The only other alternative is that some leaf contains an infinite arc (see figure). But then since this leaf is contained in S, which is compact, it must pass arbitrarily close to itself, so we can short-circuit it to a loop. But this loop, thought of in M, consists of an arc in a leaf of $\mathcal{F}$ together with a short transverse arc. But this loop can then be deformed slightly to a transverse loop (Because $\mathcal{F}$ is transversely-oriented; M has no (non-trivial) double covers!), which must of course be null-homotopic. But then the proof of Novikov's first argument goes through to show that some leaf of $\mathcal{F}$ has a vanishing cycle, which is absurd; every leaf of $\mathcal{F}$ is simply-connected, so certainly cannot contain a loop non-trivial in its fundamental group!

Now if we pick a center singularity of $\mathrm{i}^*(\mathcal{F})$ and work our way out (as in the previous argument), only two things can happen; either we bump into another center singularity (so $i^*(\mathcal{F})$ has no saddles, which is what we want), or we bump into a saddle singularity (once or twice) - the other possibilities considered before are rule out, because we can't find a vanishing cycle and S has no boundary.

If we bump into a saddle, we can (after possibly choosing a different center, as before), that we run into it only once, so we have a picture like the one below. The arc of the singular leaf that forms the loop γ we have run into is embedded in S, hence embedded in the leaf L of $\mathcal{F}$ containing it. Since it must be null-homotopic in the leaf, it therefore bounds a disk D in L. Now look at D$\cap$S; since D is contained in a leaf, this intersection, in S, consists of a (finite) collection of the leaves of $i^*(\mathcal{F})$. Since D already contains a (saddle) singularity of $i^*(\mathcal{F})$ (in its boundary), we can assume that it contains no other singularities (by having arranged that the singularities of $i^*(\mathcal{F})$ are in distinct leaves of $\mathcal{F}$). Therefore the intersection looks like, in D, a collection of simple loops, and possibly the <u>other</u> arc joining the saddle singularity to itself (see figure). If we do have the other arc, we carry out the argument below for the subdisk of D that it bounds.

The loops in this disk are are in both S and the leaf L, so they are loops of the singular foliation $i^*(\mathcal{F})$. If we pick an innermost such loop γ (see figure), then we can <u>surger</u> S along the disk D$'$ in D it bounds; we replace S with two spheres, each consisting of one of the disks in S that the loop γ bounds, together with the disk in D that γ bounds; which we then push off of one another slightly. We can deform these spheres with corners to smooth spheres which are Morse, by introducing a pair of center singularities near D$'$. Continuing by induction, we can arrange that this <u>collection</u> of 2-spheres now miss D$'$(except for the one that contains its boundary!). But then if we also surger S along D$'$ (see figure), we get a collection of 2-spheres where now we can <u>cancel</u> a center (one of the one we just created!) against our saddle (there are two cases; see figure). Therefore we get a collection of 2-spheres (obtained by surgering our original one) whose total number of saddle singularities has decreased. Then we have:

Proposition: If S$\subseteq$M is a 2-sphere and D is a disk in D with S$\cap$D=∂D=γ, splitting S into disks D^+,D^-, and both of the 2-spheres $S^+=D^+\cup D$,$S^-=D^-\cup D$ bounds a 3-ball, then S bounds a 3-ball.

Proof: The proof is by picture! If both bound 3-balls (B^+,B^- respectively), then $\partial B^+\cap\partial B^-=D$; if int($B^+$)$\cap$int($B^-$)=$\emptyset$, then it is easy to see (see figure) that $B^+\cup B^-$is a 3-ball, with boundary S; if int(B^+)$\cap$int(B^-)$\neq\emptyset$, then again, it is easy to see that one contains the other (say $B^-\subseteq B^+$), and then $B^+\setminus B^-$ is a 3-ball with boundary S.

We use this result and induction on the number of saddles in the collection of 2-spheres to show that our original sphere S bounds a ball. Since after cancelling the center and saddle, the <u>total</u> number of saddles in the 2-spheres has decreased, by induction we can assume that all of the resulting 2-spheres bound 3-balls. Then the result above (and induction on the number of 2-spheres!) implies that our original S bounds a 3-ball.

This leaves the base case - no saddles. We assume we have a 2-sphere S, in Morse normal form with no saddle singularities. Therefore the singular foliation consists of two center singularities, with a collection of (parallel) loops in between. These loops can be parametrized by a transverse arc [-1,1] running between the centers, so we will call the loops γ_t, $t\in(-1,1)$. Each is an embedded loop in a leaf L_t of $\mathcal{F}$, so bounds a disk D_t in that leaf.

Proposition: For $t_1 \neq t_2$, $D_{t1} \cap D_{t2} = \emptyset$.

Proof: If $D_{t1} \cap D_{t2} \neq \emptyset$, then since $\partial D_{t1} \cap \partial D_{t2} = \gamma_{t_1} \cap \gamma_{t_2} = \emptyset$, it must be the case that one contains the other (the alternative is that the leaf containing them is the union of two disks with disjoint boundary, which is a 2-sphere), say $D_{t1} \subseteq D_{t2}$. But then the arc $[t_1, t_2]$ in S together with an arc in D_{t2} joining γ_{t_1} to γ_{t_2} is a loop which, as usual, can be deformed to a loop everywhere transverse to $\mathcal{F}$, a contradiction.

Now we can finish the proof, using an argument similar to that which we gave in Novikov's theorem. Each disk D_{t0} can be flowed up and down the other nearby disks, and, in this case, since the disks are all embedded, this flow can be realized as an embedding of $D_{t0} \times [t_0 - \epsilon, t_0 + \epsilon]$ into M, with boundary $D_{t-\epsilon} \cup (\cup\{\gamma_t : t_0 - \epsilon \leq t \leq t_0 + \epsilon\} \cap D_{t+\epsilon}$; see figure. Notice that this is a 3-ball! Therefore, as before, using the compactness of [-1,1] we can find a finite number of D_t's so that surgering S along the entire finite collection gives a collection of 2-spheres each bounding a 3-ball $D_{ti} \times [t_i - \epsilon, t_i + \epsilon]$. But then the proposition above and induction implies that the original 2-sphere S bounds a 3-ball!

Note that we have proved, for example, that R^3 is irreducible (this is known as Alexander's Theorem), since R^3 can be foliated by (horizontal) planes. When you get right down to it, though, this is actually <u>all</u> we have proved: Palmeira has shown, in fact, that any simply-connected n-manifold that can be foliated by (n-1)-planes is homeomorphic to $\mathbf{R}^n$! But demonstrating this result would take us too far afield. We will instead continue with our current approach; our next task is to prove a generalization of the construction in this proof (centers and saddles can be cancelled against one another) for surfaces of higher genus. This result will be a key step in our proof of a theorem of Thurston that a compact leaf of a Reebless foliation is 'topologically minimal'.

Foliations and the Topology of 3-manifolds

Outline of class 15

Our next goal is to show:

Theorem (Thurston): If $\mathcal{F}$ is a (transversely-orientable) foliation without Reeb components of the (compact, orientable) 3-manifold M, and if F is a compact leaf of $\mathcal{F}$, then F has minimal genus in its homology class, i.e., for any (possibly not connected) surface S with [S]=[F] in $H_2(M)$ (or $H_2(M,\partial M)$, whichever is appropriate), we have genus(S)$\geq$genus(F).

We will need to develop two new techniques for our proof. One is a generalization of the result we used for Rosenberg's theorem - we can, by isotopy, cancel centers against saddles in the singular foliation of more general surfaces. This will be our topic for today. The second is an understanding of the Euler class of a 2-dimensional vector bundle (like $T\mathcal{F}$), and how to calculate it using a singular foliation.

The basic content of the first technique is contained in:

Theorem (Thurston, Roussarie): If $\mathcal{F}$is a foliation without Reeb components in the 3-manifold M, and $F\neq S^2$ is an embedded surface in M with $\pi_1(F)\hookrightarrow\pi_1(M)$, then F is isotopic to a surface such that the inclusion is a Morse function (except at a finite number of circle tangencies) with no center singularities in the induced singular foliation.

A surface F in M with $\pi_1(F)\hookrightarrow\pi_1(M)$ is called an (algebraically) <u>incompressible</u> surface. It is not hard to see that this condition implies the condition of (geometric) incompressibility: for any disk D in M with $D\cap F=\partial D=\gamma$, there is a disk D' in F with $\partial D'=\gamma$. This is because γ is an embedded loop in F which is null-homotopic in M, hence in F, so bounds a disk in F. It is actually true that, for 2-sided surfaces (ones for which $N(F)\setminus F$ is not connected), the opposite is true: geometric incompressibility implies algebraic incompressibility. This is the celebrated Loop Theorem. We will, however, not be making use of this direction in our proof.

Proof: The basic idea is to mimic, as much as we can, the construction we gave in our proof of Rosenberg's theorem. We start by deforming F slightly to make the inclusion a Morse function. If it has no senter singularities, we are done. Otherwise, we start at a center and start walking out, looking at the loops of the singular foliation. As before, only a few things can happen - either we run into a saddle (once or twice), we run into another center, or we run into a loop which is <u>not</u> null-homotopic in its leaf. But the last two in fact can't happen, since the first of them implies that F is a 2-sphere (ruled out by hypothesis), while the other gives us a loop in a leaf, not null-homotopic in that leaf, but demonstrably null-homotopic in M (it bounds a <u>disk</u> in F), contradicting Novikov's theorem. So we therefore run into a saddle; as before, this can happen in one of two ways have dealt with before, we cannot avoid the second case - it is why we must allow for circle tangencies.

Let us deal first with the case we already understand - when we run into a saddle once when moving out from a center. This gives us an embedded loop in a leaf L of $\mathcal{F}$; since it is null-homotopic in F (hence M - it bounds a disk in F), it is therefore null-homotopic in L, hence bounds a disk D in L. If we look at $D\cap F\subseteq D$, after possibly passing to a sub-disk as before, we can assume that this intersection consists of loops in D. If we pick an innermost such loop, bounding a disk D$'$ missing F, we get a disk satisfying the hypotheses of (geometric) incompressibility; therefore there is a disk D$''$ in F bounded by the same loop. Together these disk form an embedded sphere, which by Rosenberg's theorem bounds a ball. We can then use this ball to isotope F, taking D$''$ to D$'$; pushing it a little farther isotopes F so that (at least) this innermost circle disappears. After the isotopy F now looks like the figure below; we seem to have introduced a new center singularity on F, but because the singular foliation on D$''$ had the boundary as a leaf, there is at least one center singularity for F in the disk, so we have in fact not increased the number of centers by this isotopy.

Then continuing this process we can make D$\cap$F=∂D; then as before we can cancel the center and saddle with one another, reducing the number of centers. After doing this a finite number of times, we can assume that all of the centers that are left (if any) meet their corresponding saddles twice. Notice that we must consider this case, rather than ignore it as we have done before; there is no guarantee, since we are not in a sphere, that the two arcs that join the saddle to itself form null-homotopic loops.

We can turn two arcs containing the saddle point into a loop δ null-homotopic in its leaf L; by breaking the double point of this loop (see figure) we can make it an embedded loop γ (freely-homotopic to $delta$, hence) null-homotopic in L. Therefore it bounds a disk D in M. There are now two cases to consider, depending on which side of γ this disk is. If it contains the saddle point, then the disk pinches down to a pair of disks, each bounded by one of the arcs joining the saddle. But then inside each of these, by the construction we did before, contains a center whose saddle is of the first type, a contradiction. Therefore we can assume that the disk does <u>not</u> contain the saddle point, so this surface F intersects this disk in a finite collection of loops. Then the usual innermost loop argument given above isotopes F to a leaf missing D in its interior. Then by pushing the disk in F that this bounds to a disk slightly above D (see figure), we can assume we are in a situation in the figure; but then by pushing this disk down into the leaf L and smoothing things out, we can cancel the center and saddle, at the expense of introducing a circle tangency into the singular foliation. Doing this for all of the centers, we can remove them all by introducing circle tangencies. Therefore after an isotopy, F has no center singularities, and possibly some circle tangencies.

We should note that in many circumstances we can remove circle tangencies in F; if the loop representing the tangency has non-trivial holonomy around it in the singular foliation of F (see figure), we can, by pushing the tangency down a little to make the loop <u>transverse</u> to the singular foliation, so the inclusion is now Morse near the loop, too.

We can also improve on this result, if we assume the F is a torus and $\mathcal{F}$ is taut:

Theorem (Thurston, Roussarie): If $\mathcal{F}$ is a taut foliation of M, and T is an incompressible torus in M, then T can be isotoped so that either T is everywhere transverse to $\mathcal{F}$, or T is a leaf of $\mathcal{F}$.

Proof: By the above, we can isotope T so that the induced foliation has no centers. Since χ(T)=0, it has no saddles either. Therefore it only has circle tangencies. What we will do it to try to isotope the tangencies away. Because we can assume that the torus T 'turns around' at the tangency γ (see figure - otherwise we can pull T transverse to $\mathcal{F}$), we can try to start pushing the annulus around γ down into the nearby leaves. Now, one of three things happens: we run into another circle tangency on one side but not the other, allowing us to cancel the two circles with one another, reducing the number of tangencies and finishing our proof by induction. Or, we run into tangencies on both sides - if they are different tangencies, we can turn the three into one, and continue (actually, this can't happen - the two tangencies came from saddles, so we would have had two saddles in the same leaf), or it is the <u>same</u> circle; it this case, pushing the annulus around our first circle into the leaf containing the second, we push T into a leaf of $\mathcal{F}$, hence it <u>is</u> a leaf of $\mathcal{F}$. The final possibility is that we keep pushing our annulus down forever; although, in the torus, this looks as if we are limiting on two loops of the induced foliation of T. This situation should sound familiar; it is entirely analogous to the situation we studied when we built a Reeb component.

As we push the annulus A down into the leaves of $\mathcal{F}$, we build a function from A$\times$[0,∞) to M, and a point must trace out a trajectory that is infinitely long (otherwise, there is a limiting map of an annulus into a leaf, which we could push further). This trajectory must pile up somewhere in M, so passes through A again. Taking the first time t that the trajectory returns to A, we can create a map from A$\times$[0,t]/ , which is a torus crossed with an interval, bent at the boundary (see figure) into M. Then the points in the two annuli in the boundary which are also in leaves all have limiting values as t approaches 0. As before, these limiting leaves are then tori, and in between we see what is known as a **cylindrical component**: it looks like the analogue of a Reeb component in an annulus, crossed with a circle (see figure).

It only remains to show that a cylindrical component cannot exist in a taut foliation; this should also be a familiar argument. The leaves of a cylindrical component can be transversely oriented so that the normal vectors point everywhere into the component. But then we cannot find a loop everywhere transverse to the foliation passing through either of the torus leaves; any loop passing into the cylindrical component cannot get back out again.

It is worth noting that we <u>can</u> build Reebless foliations in, for example, the 3-torus, so that there are incompressible tori that cannot be isotoped to be transverse to the foliation or to be a leaf. Such a foliation must of course contain a cylindrical component; we leave actually constructing one (and showing that there is such an incompressible torus) as an exercise.

The above result was used by Thurston, in his thesis, to prove:

Theorem: If M is a circle bundle over a surface, and $\mathcal{F}$ is a taut foliation with no compact leaves, then $\mathcal{F}$ is isotopic to a foliation everywhere transverse to the circle fibers of M.

The basic idea of his proof was to take an essential loop in the surface; its inverse image is an incompressible torus in M, which we can then make the foliation transverse to (we imagine moving the foliation and not the torus). With a bit of work, an inductive use of the above theorem allows us to conclude his result. I've never had much motivation to look deeper into the proof, however, because of

Theorem (me!): If $\mathcal{F}$ is a foliation with no compact leaves in a Seifert-fibered space (= a (compact) manifold foliated by circles), then $\mathcal{F}$ is isotopic to a foliation everywhere transverse to the circle fibers of the Seifert-fibering.

whose proof I understand a whole lot better! The theorem is actually the subject of that reprint I had lying out on the table during this class; but its title should make it clear that the proof would take us far off the track we are now following, so we won't go through it.

Next time we will look at the Euler class of a 2-plane bundle (like T$\mathcal{F}$), and use what we learn to prove the topological minimality of leaves of Reebless foliations.

Foliations and the Topology of 3-manifolds

Outline of class 16

Last time we saw how we could isotope an incompressible surface F in a manifold M with a Reebless foliation $\mathcal{F}$ so that F had no center singularities in its induced singular foliation, at the expense of introducing circle tangencies between F and some of the leaves of $\mathcal{F}$. Therefore, all of the singularities of the foliation on F have index -1 (when viewed from F). Today we will see how to use this fact to prove Thurston's theorem; for this we will need to explore the concept of the Euler class of a bundle.

To simplify things somewhat, we will restrict our study of the Euler class to the situation in which we will actually use it. Suppose we have an orientable 2-plane bundle X_i over a reasonable space (like for example a manifold M). As an example, we have the tangent space T$\mathcal{F}$ of a transversely=-oriented codimension-1 foliation on an orientable 3-manifold M (the transverse orientation allows us to choose an orientation for the planes of our bundle, by making it the first element of an orthonormal frame at a point; the other two vectors, in order, give us a frame for the tangent plane. Then the Euler class e(X_i) of X_i is an element of H^2(M), the second cohomology of M with coefficients in **Z**, which represents the obstruction to the existence of a nowhere-zero section of X_i (e.g., a nowhere-zero tangent vector field, in our case).

To define it, we will define a 2-cochain (a function from formal sums of 2-dimensional simplices to **Z**) which is zero on the boundaries of 3-simplices (hence is a 2-cocycle). To make it the obstruction to the existence of a nowhere-zero section, we define it in terms a partially-constructed nowhere-zero sections.

Given a 2-simplex in M σ^2, we can arbitrarily choose non-zero vectors at its vertices. Then, since the 1-simplices in its boundary are contractible (so the bundle X_i, restricted to the 1-simplex, is trivial), we can define a nowhere-zero vector field over each 1-simplex, and we can (by changing it in the vicinity of the vertices, as necessary) assume it agrees with the 'vector field' already defined at the vertices (see figure). But since the 2-simplex itself is contractible, the bundle $X_i|_{\sigma^2}$ is itself trivial, so we get a commutative diagram like below, which allows us to define a map from $\partial\sigma^2$ to S^1. This map is a map from an (oriented - 2-simplices carry a standard 'counterclockwise' orientation) circle to an (oriented - the circle inherits an orientation from the orientation of the 2-planes (again, counterclockwise)) circle, so it has a well-defined winding number w(σ^2). Then we define e(X_i) of a 2-chain $\Sigma\sigma_\mathrm{i}$ to be e(X_i($\Sigma\sigma_\mathrm{i}$)=Σw(σ_i).

However, this function involves alot of choices, which we must show that it is in some sense independent of. In particular, this function was defined in terms of a choice of nowhere-zero section over the 1-simplices in the boundary of our 2-simplices. Clearly, changing this section (by, for example, adding several windings along one of the 1-simplices - see figure) will certainly change the winding numbers, so will change the value of the Euler class. The point, however, is that for 2-cycles, the value of the sum is independent of these choices. This is because if we change our choice of vector field on a 1-simplex, it changes the winding numbers associated to every 2-simplex that contains it. So since for a 2-cycle, which is a 2-chain for which the numbers associated to each of its boundary 1-simplices (with signs coming from orientations) sum to zero, any changes we make to the sections cancel out, so the value the Euler class assigns is unchanged (see figure below)

This at least gives us a function from 2-cycles to **Z**. With a little bit more work, you can actually show that e(X_i) is a well-defined function from 2-chains to **Z**, up to the choice of a 2-coboundary, i.e., the difference of functions defined by two choices of sections over 1-simplices is a 2-coboundary. To show that e(X_i) os a 2-cocycle, it suffices only to show that it takes value 0 on any 2-boundary, i.e., it is 0 on the sum of 2-simplices in the boundary of a 3-simplex. But since we've seen that how we define the section on the 1-simplices is irrelevant, we can simply choose a section that makes it obvious that the resulting sum is 0 (since all of the winding numbers are 0 - see figure). So we get an element e(X_i)∈.

From this definition it is easy to see that e(X_i) is natural; if F is a subcomplex of M (e.g., an embedded surface in M), then e$X_i|_F$=i^*(e(X_i)∈H^2(F), where i^* is the map on cohomology induced from the inclusion map i. This is because the winding number around the boundary of a 2-simplex will be the <u>same</u>, whether we pretend the simplex is in F or is in M. (In fact, this naturality extends to any continuous map from a space X to M; the Euler class pulls back to the Euler class of the pull-back 2-plane bundle. In other words, the Euler class of the pullback is the pullback of the Euler class!)

It is also easy that the Euler class of a bundle over an (orientable) <u>surface</u> F can be calculated by counting indices of a singular vector field on F. To see this, just remember that H_2(F)=**Z**, generated by the fundamental class [F] (which is just the sum of all of the simplices of F, for some triangulation), so e(X_i)∈H^2(F)=**Z** is basically determined by its value on [F]. But if we choose a singular vector field on F with at most one singularity in each 2-simplex (think of this as starting with a singular vector field and choosing a fine enough triangulation around it). then the winding numbers of simplices which have no singularity in it is zero, since winding number is really the obstruction to having a map $S^1 \to S^1$ which <u>extends</u> to $D_2 \to S^1$, and simplices without a singularity clearly extend their maps on the boundary. But then the Euler class evaluated on [F] is then the winding numbers of the vector fields around the remaining simplices, which are disks each containing a simplex, and this is basically the index of the vector field at that point (where some care must be made in determining the sign, since that is a matter of convention between the orientation of the simplex (which orients its boundary) and the orientation of the bundle (which orients the fiber we projected to)).

So for example, of F is an (orientable) surface and X_i=TF is its tangent bundle, then e(TF)([F]) can be calculated by finding a vector field on F with isolated zeros (which is the short way of saying a vector field in TF), and adding the indices of zeros together, keeping track of the sign conventions. But in this case the orientation on TF <u>comes</u> from the one on F, so the orientations are so chosen so that the index is equal to the winding number. Consequently, e(TF)([F])=the sum of the indices of the zeros of a vector field on F. But we've already seen how to give a different name to this number; it is the Euler characteristic of F. So we have shown:

$$e(TF)([F])=\chi(F) \text{ for any orientable surface.}$$

We are now just about ready to apply these ideas to prove Thurston's theorem. What we have is a foliation $\mathcal{F}$ without Reeb components of a 3-manifold M, and a compact leaf F of $\mathcal{F}$. We also have a (possible disconnected) surface S representing the same homology class as F in M: [S]=[F]∈H_2(M). Therefore we of course have:

$$e(T\mathcal{F}|_S)[S] = e(T\mathcal{F})[S] = e(T\mathcal{F})[F] = e(T\mathcal{F}|_F)[F]$$

where the middle equality is because [S]=[F], and the outer two equalities are because of the naturality of the Euler class. But because F is a leaf of $\mathcal{F}$, $T\mathcal{F}|_F$=TF, so $e(T\mathcal{F}|_F)[F] = e(TF)[F] = \chi(F)$=2-2genus(F). Consequently, we have:

$$e(T\mathcal{F}|_S)[S] = 2 - 2\text{genus}(F)$$

Therefore, if we can show that

$$\chi(T) \leq e(T\mathcal{F}|_T)[T]$$

for every component T of S, then we would be done: we would then have

2-Σ2genus(T)$\leq\Sigma$(2-2genus(T))=$\Sigma(\chi(T)) = \Sigma(e(T\mathcal{F}|_T)[T]) = e(T\mathcal{F}|_S)[S]$=2-2genus(F), so Σ(genus(T))=genus(S)$\geq$genus(F).

This is what we will in fact show, at least for T not equal to a 2-sphere. Notice that technically the statement is false for 2-spheres, since Rosenberg's theorem tells us that any embedded 2-sphere bounds a 3-ball, hence is null-homologous (it's a boundary), so $e(T\mathcal{F})$ evaluates to 0, which is smaller than $\chi(S^2)$=2. But since 2-spheres add nothing to the homology class [S], we can throw them away without affecting [S]=[F]; we can technically add the statement that no component of S is a sphere to Thurston's theorem, without reducing any of it's scope, and so we do this.

We will do this by building a singular foliation on T and applying the work we have done above to understand singular foliations on T and how they relate to $e(T\mathcal{F}|_T)[T]$. But in order to apply the isotopy theorem of last time, we actually need an incompressible surface; we will show how to arrange this part, and leave the rest of the proof, which calculates $e(T\mathcal{F}|_T)[T]$, for next time. The idea is that if we don't have an incompressible surface T, then there is a disk D in M with $D\cap T=\partial D\subseteq T$, as we argued before. Then if we surger T along D we get a new (possibly disconnected) surface T′, with $\chi(T')=\chi(T)+2$. But [T′]=[T] in $H_2(M)$, since together the two surfaces are the boundary of a 3-dimensional piece of M (see figure). Therefore, if we know $\chi(T')\leq e(TF)[T']$, then we have $\chi(T)\leq \chi(T')\leq e(TF)[T']=e(TF)[T]$, and we would be done (by induction, basically). But we know that every time we surger T either (we surger along a non-separating curve and) the genus of T goes down by one or (we surger along a separating curve that does not bound a disk on either side, so) T splits into two pieces each with smaller genus than T. So the process cannot continue forever and the eventually end with either a bunch of spheres (whose sum is null-homologous so [T]=0 so $\chi(T)\leq=e(T\mathcal{F}|_T)[T]$) or a bunch of spheres and incompressible surfaces, so T is homologous to a union of incompressible surfaces. If we verify our basic inequality for each of the incompressible surfaces T_i, then $\chi(T)\leq \Sigma\chi(T_i)\leq\Sigma(e(TF)[T_i]=e(TF)[T]$, and we would be done.

Therefore we need only prove the basic inequality for incompressible surfaces T in M; for this we will be able to apply the isotopy theorem we proved last time, and the counting techniques we developed this time, to furnish a proof. This will be done next time.

Foliations and the Topology of 3-manifolds

Outline of class 17

This time we will finish the proof of Thurston's theorem.

We have so far reduced the proof to showing that, for any incompressible (orientable) surface S in M, $\chi(S) \leq e(T\mathcal{F}|_S)[S]$. We can therefore isotope S, by the Thurston-Roussarie result, so that the induced singular foliation $i^*\mathcal{F}$ on S has no center singularities, only saddles and circle tangencies. Notice that the isotopy does not change either side of the inequality we wish to prove - we still have the same surface, and the two are homologous (the track of ther isotopy provides a 3-boundary between them). We will use the singular foliation to calculate the two quantities, by building vector fields in T(S) and $T\mathcal{F}|_S$ whose zeros correspond exactly to the singularities of the foliation on S. Then calculating each side, using indices and winding numbers around these singularities, respectively, will finish off the proof.

Suppose first that the singular foliation on S has no circle tangencies. Then it turns out that we can build a <u>single</u> vector field to suit our purposes - it can be thought of alternately as living on S (i.e., living in T(S)) and living in $T\mathcal{F}|_S$. To build it, note that away from the saddle singularities (where the leaves of $\mathcal{F}$ are tangent to S), the leaves of $\mathcal{F}$ are <u>transverse</u> to S (see figure). Therefore at each point of S (other than the saddles) there is a well -defined choice of normal vector in S to the leaf of the singular foliation which lies within 90 degrees of the normal to the leaf of the foliation (given by a transverse orientation) at that point. For otherwise, <u>both</u> normals are orthogonal to the vector orthogonal to the leaf, so are both <u>tangent</u> to the leaf, implying the leaf is tangent to S at that point, a contradiction. These vectors form a vector field on S (the choice of normal is locally constant), except at the saddles, which is orthogonal to the leaves of the singular foliation. If we now rotate all of these vectors 90 degrees to the right (w.r.t. some orientation of S), we get a (singular) vector field on S which is tangent to the singular foliation $i^*(\mathcal{F})$.

Notice that this vector field is tangent to S (by construction), but because it is tangent to the leaves of the singular foliation, it can <u>also</u> be thought of as a vector field i**n** $\mathbf{T(\mathcal{F}|_S)}$. This is basically because locally it looks like the picture below. The <u>point</u> is that this one vector field can therefore be used to calculate **both** $\chi(S)$ **and** $e(T\mathcal{F}|_S)[S]$. Before doing so, though, we need a definition:

Definition: A saddle singularity p of $i^*(\mathcal{F})$ is called **positive** if the normal vector to S at p (chosen using the orientations of M and S, so that an orientation of S, followed by the normal, gives an orientation of M) is the <u>same</u> as the normal to the leaf of $\mathcal{F}$ at p (coming from the transverse orientation of $\mathcal{F}$) - see figure. If they are opposite, instead, the saddle is called **negative**.

Let I_P = the number of positive saddles of $i^*(\mathcal{F})$, and let I_N = the number of negative saddles.

Now, calculating $\chi(S)$ using our vector field is easy. It is the sum of the indices of its zeros, all of which are saddles of the singular foliation $i^*(\mathcal{F})$. So all of the indices are -1, so $\chi(S)$ is -1 times the number of saddles, which is $I_N + I_P$. So $\chi(S)=-(I_N + I_P)$.

For $e(T\mathcal{F}|_S)[S]$, the calculation is almost as easy. It is the sum of the winding numbers of our vector field around the singularities, which is calculated in almost the same way, except that a possible difference in orientations must be taken into account (this does not arise for $\chi(S)$, because the orientation for S at a point is the <u>same</u> as the orientation of the tangent space at that point (more or less by definition)). At a positive saddle (see figure), the orientations for the surface and the tangent plane agree (since leaves are oriented so that after adding the transverse orientation, we get the orientation of M), so that the winding number of the vector field around that point is -1; but at a negative saddle, the orientation of the tangent plane is <u>opposite</u> (taking an S-centered view), which forces us to count our winding number in the opposite direction, so the winding number is 1. So $e(T\mathcal{F}|_S)[S]$, which is the sum of the winding numbers around the singularities is a sum of **1's** and -1's, depending on whether the saddle is positve or negative, so we conclude that $e(T\mathcal{F}|_S)[S] = I_N - I_P$. But since $I_N \geq 0$, it is clear that $\chi(S) = -(I_N + I_P) = -I_N - I_P \leq I_N - I_P = e(T\mathcal{F}|_S)[S]$, as desired!

This finishes the case that $i^*(\mathcal{F})$ has no circle tangencies. If there are circle tangencies, we will show that we can still build <u>two</u> vector fields, one in T(S), the other in $T\mathcal{F}|_S$, which each agree with the one described above outside of small neighborhoods of the circle tangencies. They also have the further property that their singularities are also exactly the singularities of $i^*(\mathcal{F})$, so the calculations we did above go through without any change to give the same result - no new singularities means that the index/winding number calculations are concentrated where the two vector fields agree!

To do this, look at the neighborhood of a circle tangency, and the way that the procedure above would build the vector field - see below. Our problem is that that recipe breaks down at the circle tangency γ - both normals in S to the tangent circle are orthogonal to the normal to the leaf of $\mathcal{F}$ containing the circle, so there is no coherent way to choose one or the other. But none of the nearby loops (we can assume there is no holonomy around the tangent circle by an argument given before) - have this problem, so the procedure above <u>does</u> create a vector field tangent to them. This vector field looks like the one in the figure - they go around γ in one direction on one side, and in the opposite direction on the other. So there is no coherent way to extend it to γ. What we do instead is to alter the vector field near γ - but we must do it in two different ways, depending upon which 2-plane bundle we want the vector field to live in.

If we wish to create a vector field in T(S), this is relatively easy. We just alter the vector field between two loops on either side of γ as in the figure below. Basically, we are replacing the singular foliation $i^*(\mathcal{F})$ in between the loops to a 'Reeb annulus' like the one below, and taking a vector field tangent to it. This gives us a singular vector field on S, agreeing with the one produced above (away from the circle tangencies), whose zeros correspond exactly to the saddles of $i^*(\mathcal{F})$.

To build the vector field in $T\mathcal{F}|_S$, we need only be slightly more clever. The local picture is as below, with the vector field as above away from the circle tangency γ. This time we want to extend the vector field over the same little annulus around γ, except this time we want to make sure our vector are always tangent to the leaves of $\mathcal{F}$, instead of tangent to S. With a little thought, it's not too hard too see how to do this; in the local picture we just want to keep choosing <u>horizontal</u> vectors which go from pointing left to pointing right b the time we cross the annulus. We can do this just by rotating the vectors starting at one side, so that we complete a half-turn as we cross the annulus. Just choose a consistent direction to turn (like always counterclockwise) as you go across (the figure below gives a view seen from looking straight down onto the annulus from above). This gives us a vector field in $T\mathcal{F}|_S$ which agrees with the one described by rotating the normal one (away from the circle tangencies) and therefore agrees with the one built in the previous paragraph, and whose zeros agree exactly with that vector field. So we can employ the argument given before to give the inequality we seek, completing our proof of Thurston's theorem.

We finish with a few comments about this theorem:

(1): This theorem has an exact analogue for foliations of manifolds with boundary. One assumes that the foliation $\mathcal{F}$ is everywhere transverse to the boundary (usually), and that in addition to having no Reeb components, the foliation has no 'half-Reeb components' (obtained by cutting a Reeb solid torus down the middle by an annulus, transverse to the foliation, containing the core circle). Much of what we have done goes through unchanged (or with at most minor adjustments) to yield a 'relative' version of this theorem: compact leaves of such Reebless foliations have least genus in their homology class, where this last part must be interpreted as occurring in $H_2(M,\partial M)$.

(2): Thurston's theorem has a converse:

Theorem (Gabai): If a compact orientable surface $F\subseteq M$ (properly embedded, if $\partial F\neq \emptyset$) has smallest genus among all surface representing the same homology class (in

$H_2(M,\partial M)$ or $H_2(M)$, whichever is appropriate), then there is a Reebless foliation $\mathcal{F}$ for which F is a leaf.

We will not prove this theorem; but we will explore several of the concepts that go into its proof, so that we can better understand how such a foliation could be built. We will start by looking at what is known as the Thurston norm on $H_2(M)$ and $H_2(M,\partial M)$, which provides a language used heavily in Gabai's proof. It is defined more or less to be the least genus of a surface representing the homology class, which is a pretty straightforward notion, but it turns out to have some rather surprising properties.

Foliations and the Topology of 3-manifolds

Outline of class 18

Today we will begin discussing one of the largest elements of Gabai's proof of the converse to Thurston's theorem. This is Thurston's (semi-)norm on the homology of a 3-manifold. But before we define it and begin exploring some of its properties, we need to take a crash course in homology and cohomology, from a low-dimensional topologist's point of view.

Th e first fact that we will use is that cohomology is (what is known as)a **representable functor**. This means that for any n, and any abelian group G, there is a space A the n-dimensional cohomology of a space X with coefficients in G, $H^n(X;G)$, is in a (natural) one-to-one correspondence with the set [X,A] of homotopy classes of maps $f:X \to A$. The space A is known as an Eilenberg-MacLane space $A=K(G,n)$, and is characterized by the fact that its homotopy groups $\pi_k(A)=0$ for $k\neq n$, and $\pi_n(A)=G$. The correspondence, in the direction $F:[X,A]\to H^n(X;G)$, is not too hard to write down, using some standard facts about Eilenberg-MacLane spaces and homology. Given an element $[f]\in[X,A]$, it induces a map $f^*:H^n(A;G)\to H^n(X;G)$ (which depends only on [f] - homotopic maps give the same map on cohomology). But by the universal coefficients theorem, $H^n(A;G)=\mathrm{Hom}(H_n(A),G)\bigoplus\mathrm{Ext}(H_{n-1}(A),G)$ (whatever that all means). But for Eilenberg-MacLane spaces, $H_n(A)=\pi_n(A)=G$, and $H_{n-1}(A)=\{0\}$, so $H^n(A;G)=\mathrm{Hom}(G,G)$ (the Ext term is 0 since (whatever it is) it vanishes for free abelian groups). But Hom(G,G) has a canonical element, namely the identity map I; the correspondence then is given by $[f]\leftrightarrow f^*(I)$.

In particular, for n=1 and $G=\mathbf{Z}$ (the case we will be interested in), we have

$$H^1(X;\mathbf{Z})=H^1(X)=[X,K(\mathbf{Z},1)].$$

But a $K(\mathbf{Z},1)$ is easy to find, namely S^1; $\pi_1(S^1)=\mathbf{Z}$, and since its universal cover is (R^1 hence) contractible, all higher homotopy groups (agree with those of R^1 hence) are trivial. In particular, if X=M is an (orientable, compact) n-manifold, then Poincaré duality says that $H^1(M) \cong H_{n-1}(M, \partial M)$, and the correspondence $[M,S^1] \rightarrow H_{n-1}(M, \partial M)$ has a much more useful description. Given the (class of) maps [f], we can take a representative $f:M \rightarrow S^1$ and deform it to be transverse to a point $x \in S^1$. Then transversality says that $f^{-1}(x)=N$ is a compact, transversely-orientable (hence orientable) compact (n-1)-manifold (with boundary, probably), and therefore represents a relative (n-1)-cycle in $(M,\partial M)$, so gives an element of $H_{n-1}(M, \partial M)$. It takes some work to show that this homology class is independent of the choices that go into its construction; we will skip this. It is also true that relative cohomology classes can be represented by maps; for example, $H^1(M, \partial M)$ can be identified with homotopy classes of maps of pairs $[(M,\partial M),(I,\partial I)]$, where $I=[0,1]$. Using a similar argument to the above we can therefore represent elements of $H_{n-1}(M)=H^1(M, \partial M)$ by embedded (n-1)-manifolds (which miss the boundary).

But most of this is not really relevant to what we want to accomplish, except that it tells us (restricting attention now to M = an orientable, compact 3-manifold) that elements of $H_2(M, \partial M)$ and $H_2(M)$ can always be represented by a 2-cycle which is in fact an embedded compact surface.

Now from the point of view of Reebless foliations, the surfaces representing homology classes that are of the most interest are the least genus ones (provided we ignore 2-spheres). This leads us very naturally to make the following sequence of definitions.

Definition: If S is a connected, orientable, compact surface then we define $\chi_-(S)$ to be equal to 0 if S is a disk or sphere, and equal to $-\chi(S)$, otherwise. For T a (possibly non-connected) orientable compact surface, we define $\chi_-(T)=\Sigma\chi_-(S)$, where we sum over all of the components of T.

Basically, it's the Euler characteristic of T, after we throw out all of the pieces with positive χ, then change sign to make it non-negative. This should remind us of what we did in Thurston's theorem - we showed that after ignoring spheres, a leaf of a Reebless foliation had lower Euler characteristic than any other representative of the same homology class; in this terminology, it had minimal χ_-. (What we proved was actually technically stronger; we also proved that the leaf had minimal total genus - one can have two collections of surfaces with the same χ_- but different total genera.) With this in mind, it then seems natural to take the next step:

Definition: For any compact, orientable 3-manifold M, we define the (Thurston) (semi-)norm on $H_2(M)$ to be the function $x:H_2(M) \rightarrow \mathbf{Z}$ defined by

$$x(\alpha) = \min\{\chi_-(F) : F \text{ is a surface with } [F]=\alpha \}.$$

We make the obvious analogous definition for $H_2(M, \partial M)$

Of course, we can't just call a function a norm; we must actually show that it is. This is contained in the following proposition:

Proposition: For any $\alpha,\beta\in H_2(M)$ (resp., $H_2(M, \partial M)$), and any $n\in\mathbf{Z}$, we have

$$\text{(a): } x(\alpha+\beta)\leq x(\alpha) + x(\beta)$$
$$\text{(b): } x(n\alpha) = |n|x(\alpha)$$

Proof: (a): Choose surfaces S and T with $[S]=\alpha$, $[T]=\beta$, and $\chi_-(S)=x(\alpha)$, $\chi_-(T)=x(\beta)$ (this is possible since χ_- takes positive integer values; the minimum is achieved). We will show how to construct a surface R with $[R]=\alpha + \beta$, and $\chi_-(R) = \chi_-(S) + \chi_-(T)$; since $x(\alpha + \beta) \leq \chi_-(R)$ (it's a minimum), this will prove (a).

Isotope S so that it intersects T transversely; then $S\cap T$ consists of a finite number of circles (and arcs, if we are dealing with the $H_2(M, \partial M)$ case). If we think of these as lying in S, then some number of them will bound disks in S (or cut off disks from S) - see the figure below. But if we choose an innermost such loop (or outermost such arc) and then <u>surger</u> T along the disk (whose interior misses T) that this provides (see figure), then we can create a new surface T'. which is homologous to T (we've seen this before), and so is another representative of the homology class β.

$\chi(T') = \chi(T) + 2$ or 1 (depending one whether we surgered along a circle or arc - this is easy to check), so we would expect χ_- to have gone down. This can't happen ($\chi_-(T)$ is minimal, so it must be that we have created new sphere or disk components by this process. But then it is easy to check that, in all cases (in the case of a loop, we created one extra sphere, we cut one sphere into two, or we created one extra disk (and therefore two, since we must have cut open an annulus), or, in the case of an arc, we split off a disk, or cut one disk into two - basically, when we split off a disk or sphere, ignoring that piece (except for the annulus case) we get a surface <u>homeomorphic</u> to the one we started with) χ_- remains unchanged. Doing this for all trivial loops and arcs of intersection in S, and then by symmetry all trivial loops and arcs of intersection in T, we can assume we have minimal χ_- surfaces with no trivial loops or arcs of intersection.

Then we preform a cut and paste on all remaining curves of intersection, using the transverse orientations of the surfaces as a guide (see figure), we get a new surface, denoted S   T, but which we will denote R, because I lied, and TeX does <u>not</u> have a control sequence to denote this character. Now, $[R] = [S] + [T]$ in $H_2(M)$ (resp., $H_2(M, \partial M)$); a homology can be seen between them in the figure below. Also, $\chi(R) = \chi(S) + \chi(T)$, since obtaining R from $S\cup T$ basically amounted to removing two copies of each arc or loop of $S\cap T$, one from each surface, and replace it with two more, sewn in differently (think of this on the level of 1-cells, and count $\chi(R)$ combinatorially), so the Euler characteristic remains unchanged. But it is also true that (*) $\chi_-(R) = \chi_-(S) + \chi_-(T)$, since each of these is basically calculated by first throwing away any spheres or disks and then computing the Euler characteristic. But any sphere or disk in S or T survives (without any cutting and pasting) to R, since if, e.g., a sphere or disk of S intersects T, each intersection would be a trivial loop or arc of $S \cap T$, which don't exist. On the other hand, every sphere or disk of R comes from either S or T, because otherwise it was pieced together from parts of $(S \cup T) \setminus (S \cap T)$ (see figure); but one of those pieces must be an outermost disk or innermost disk, contradicting our preliminary surgery process. So the same collection of spheres and disks are thrown away on each side of the equality (*), giving an equality of χ_-'s, and finishing our proof.

Next time we will handle part (b) of the proposition, and continue to explore some of the properties of this metric. The next step is to show that these norms extend, in an essentially unique way, to continuous norms for $H_2(M;\mathbf{R})$ and $H_2(M,\partial M;\mathbf{R})$, taking values in $\mathbf{R}$.

Foliations and the Topology of 3-manifolds

Outline of class 19

Last time we introduced the Thurston norm x on second (absolute and relative) homology of a compact orientable 3-manifold, using the fact that homology classes can be represented as embedded surfaces, and defined to be, basically, the least genus of such a surface. This function turns out to be a semi-norm; this is contained in

Proposition: For any $\alpha,\beta \in H_2(M)$ (resp., $H_2(M,\partial M)$), and any $n \in \mathbf{Z}$, we have

(a): $x(\alpha+\beta) \leq x(\alpha) + x(\beta)$
(b): $x(n\alpha) = |n| x(\alpha)$

The first part was proven last time; today we prove the second, and then show how x can be extended to a semi-norm on the real second homology groups.

To prove (b), we will show that each side is smaller than the other. Without loss of generality, we may assume that $n>0$, This is because if $n=0$ the statement is obvious ($x(0)=0$); and if $n<0$, then since $x(-\alpha)=x(\alpha)$ (they are represented by exactly the same surfaces, but with opposite orientations, and χ_- ignores orientation),

$$x(n\alpha)=x((-n)(-\alpha))=-nx(-\alpha)=|n|x(\alpha)$$

(provided we have proven it for $n>0$).

Showing $x(n\alpha) \leq nx(\alpha)$ is not hard; If we pick a surface F with $[F]=\alpha$ and $\chi_-(F)=x(\alpha)$, then n parallel copies of F (call it nF) represents $n\alpha$. But then $\chi_-(nF)=n\chi_-(F)=nx(\alpha)$, and since $[nF]=n\alpha$, $x(n\alpha) \leq \chi_-(nF)=nx(\alpha)$.

For the other inequality, choose a surface T with $[T]=n\alpha$, and $\chi_-(T)=x(n\alpha)$. We will show:

Claim: T is the (disjoint) union of n surfaces $S_1, \ldots, S_n$ each representing α.

This will prove (b), since then

$$x(n\alpha)=\chi_-(T) = \chi_-(S_1\cup\cdots\cup S_n)=\Sigma_{i=1}^n\chi_-(S_i) \geq \Sigma_{i=1}^n x(\alpha)=nx(\alpha)$$

(because $x(\alpha)\leq\chi_-(S_i)$, since $[S_i]=\alpha$).

To prove the claim, we represent nα as a map f from M to S^1, with $f^{-1}(pt.)=T$. We can do this directly; T is an orientable surface, so a neighborhood of it, N(T), looks like T$\times$I, if we then map M to S^1 by sending N(T) to I (by projection) and sending I to a loop running once around S^1, and then sending M$\backslash$N(T) to the basepoint of this loop, this is a continuous map, and T is exactly the inverse image of a point (on the far side of S^1 from the basepoint). α can alse be represented as a map g:M$\rightarrow S^1$, and if we then compose g with the standard n-fold covering map p:$S^1\rightarrow S^1$, we also get a map p$\circ$g:M$\rightarrow S^1$ representing nα. But since there is a correspondence between (cohomology, hence) homology classes and homotopy classes of maps from M yo S^1, it must therefore be the case that f and p$\circ$g are homotopic to one another. But since p is a covering map, the Covering Homotopy Theorem implies that there is a map g$'$:M$\rightarrow S^1$ homotopic to g (hence representing the same homology class) such that p$\circ$g$'$=f. But then $T=f^{-1}(*)=(p\circ g')^{-1}(*)=(g')^{-1}(p^{-1}(*))= (g')^{-1}(*_1\cup\cdots\cup *_n)= (g')^{-1}(*_1)\cup\cdots\cup(g')^{-1}(*_n)= S_1\cup\cdots\cup S_n$, where each surface S_i is the inverse image of a point under g$'$, so represents α, as desired.

This finishes the proposition, and shows that x is a semi-norm. It is not a norm in general: it would have to satisfy the additional requirement that $x(\alpha)=0$ implies $\alpha=0$.. It is easy to see exactly when this is true: x is a norm if no non-trivial homology class can be represented by a union of spheres and tori (and disks and annuli, in the relative case), since these are the only surfaces with $\chi_-=0$.

Dealing with this semi-norm on the level of homology with integer coefficients is fine, but ends up obscuring much of the structure that this semi-norm has. Now, $H_2(M)$ and $H_2(M,\partial M;\mathbf{R})$ are torsion-free (as is the (n-1)th-dimensional homology of any n-manifold). So we can think of it as a 'vector space' $\mathbf{Z^n}$ over $\mathbf{Z}$. Consequently, by the universal coefficients theorem, the second homology with coefficients in a field is an n-dimensional vector space over the field, and we can think of it as having the 'same' basis as the $\mathbf{Z}$-vector space. In particular, we can imagine $H_2(M)$ sitting in $H_2(M;\mathbf{R})$ as its integer lattice (and similarly for relative homology). What we will do now is extend x to a semi-norm on homology with real coefficients.

Proposition: x extends uniquely to a continuous semi-norm $x:H_2(M;\mathbf{R})\rightarrow\mathbf{R}^+$ (resp. $H_2(M,\partial M;\mathbf{R})$).

Proof: We start by extending x to $\mathbf{Q^n}\subseteq\mathbf{R^n} =H_2(M;\mathbf{R})$ (or $H_2(M,\partial M;\mathbf{R})$, as desired). We want to extend it in such a way that it remains convex and linear on rays - the point is that these two criteria force us to extend it in only one way.

Given any $\alpha\in\mathbf{Q^n}$, some multiple of it $k\alpha\in\mathbf{Z^n}$, so $x(k\alpha)$ is defined. If we are to define $x(\alpha)$ so that we have linearity on rays, then we must have $|k|x(\alpha)=x(k\alpha)$. But we can then use this equation to define $x(\alpha)$; just set $x(\alpha)=|k^{-1}|x(k\alpha)$. But this definition involves a choice of k, so we need to make sure that the result is independent of k. This is not hard; if $k\alpha, m\alpha \in \mathbf{Z^n}$, then $km\alpha \in \mathbf{Z^n}$, and by the linearity of x on $\mathbf{Z^n}$, we have $x(km\alpha)=|k|x(m\alpha)=|m|x(k\alpha)$, so dividing the last two by $|km|$, we get $|m^{-1}|x(m\alpha)=|k^{-1}|$ $x(k\alpha)$, so the definition of $x(\alpha)$ is independent of k.

It is easy to see that the two properties for a semi-norm are true for this extension to $\mathbf{Q^n}$ Linearity follows from the equalities $x(n\alpha)= |k^{-1}|x(kn\alpha)=|k^{-1}||n|x(k\alpha)=|n|x(\alpha)$ (by choosing a suitable k). Convexity follows by choosing a k so that $k\alpha$ and $k\beta$ are both in $\mathbf{Z^n}$, and then $x(\alpha + \beta) = |k^{-1}|x(k(\alpha + \beta)) \leq |k^{-1}|x(k\alpha) + |k^{-1}|x(k\beta) = x(\alpha + x(\beta)$.

Convexity implies that this function $x\mathbf{Q^n} \rightarrow \mathbf{Q^+}$ is continuous. This is because $x(\alpha + \beta) \leq x(\alpha + x(\beta)$, the triangle inequality for the norm, implies that $|x(\alpha)-x(\beta)| \leq x(\alpha-\beta)$ (this is standard). Therefore to show that x is continuous, it suffices to show that if α is small, then $x(\alpha)$ is small. But if we write $\alpha=(a_1,\ldots a_n)$, then $x(\alpha) \leq |a_1|x(1,0,\ldots,0) + \cdots + |a_n|x(0,\ldots,0,1) \leq (a_1 + \cdots + a_n)C$, where $C=\max\{x(1,0,\ldots,0),\ldots,x(0,\ldots,0,1)\}$. But if α is small, then each of the a_i are, so their sum is, so C times the sum is. So x is continuous.

Now we extend this extesion to $\mathbf{R^n}$. Since $\mathbf{Q^n}$ is dense in $\mathbf{R^n}$, given any $\alpha \in \mathbf{R^n}$, we can find $\alpha_n \in \mathbf{Q^n}$ converging to α, so in particular the α_n form a Cauchy sequence in $\mathbf{Q^n}$, so they are close to one another (i.e., their differences are small). But then the calculation above then implies that the $x(\alpha_n)$ are close to one another, i.e., they form a Cauchy sequence in $\mathbf{R}$, and hence converge to some number, which we define to be $x(\alpha)$. This definition is independent of the choice of sequence, because if we choose another sequence β_n converging to α, then $\alpha_n-\beta_n$ converges to 0, so by the above, $x(\alpha_n-\beta_n)$ converges to 0, so $x(\alpha_n)-x(\beta_n)$ converges to 0, i.e., the two sequences converge to the same number. This extension to $\mathbf{R^n}$ can be easily seen to satisfy the two criteria for a semi-norm; if we choose α_n converging to α, choose $k\alpha_n$ converging to $k\alpha$, so $x(k\alpha_n)=|k|x(\alpha_n)$ converges both to $x(k\alpha)$ and $|k|x(\alpha)$, so they are equal. For convexity, since $x(\alpha_n+\beta_n) \leq x(\alpha_n)+x(\beta_n)$, and the left side converges to $x(\alpha+\beta)$ while the right side converges to $x(\alpha)+x(\beta)$, the first must be less than or equal to the second.

Finally, this extension to $\mathbf{R^n}$ is continuous, by the same argument as given above, since we have shown that it is convex. Note that this last extension is also forced on us, if we want to have continuity, since that requires that convergent sequences be carried to convergent sequences.

It is easy to see that the set of points of norm 0 is a linear subspace of $\mathbf{R^n}$; if $x(\alpha)=x(\beta)=0$, then $x(a\alpha+b\beta) \leq |a|x(\alpha)+|b|x(\beta)=0+0=0$, so since it is positive, $x(a\alpha+b\beta)=0$. $\blacksquare$ If the function x originally defined on $\mathbf{Z^n}$, is a norm, however, then it is easy to see that this extension is a norm (meaning $x(\alpha)=0$ iff $\alpha=0$). The converse is also true; but since nobody seems to care about that, we won't prove it here.

We can in any case define a unit ball for the norm x; $B_x=\{\alpha\in\mathbf{R^n} : x(\alpha)\leq 1\}$. However, this ball is most interesting when x is a norm. We start with

Proposition: If x is a norm, then B_x is compact.

Proof: We're in $\mathbf{R^n}$, so it's enough to show that B_x are closed and bounded. Closed is easy: $B_x= x^{-1}([0,1])$, which is closed since x is continuous. for bounded, we work projectively: if not, then there is a sequence α_n converging to ∞ with $x(\alpha_n)\leq 1$ for all n, and so $\alpha_n/||\alpha_n||$ lies on the unit sphere in $\mathbf{R^n}$ and $x(\alpha_n/||\alpha_n||)\leq 1/||\alpha_n||$ converging to 0. But this gives a contradiction, since the sequence in the sphere has a convergent subsequence (the sphere is compact), and by continuity of x it must converge to a point α with $x(\alpha)=0$, which since x is a norm, is therefore 0, which doesn't lie in the sphere!

If B_x is compact (which if you think about it, implies the x is a norm), then we can use this to define a norm x^* on the dual space $\mathrm{Hom}(H_2(M;\mathbf{R}),\mathbf{R})=H^2(M;\mathbf{R})= H_1(M,\partial M;\mathbf{R})$ (equalities are by the various theorems of homology theory) by

$$x^*(L)=\max\{|L(\alpha)| : x(\alpha)\leq 1\}$$

This is finite, for any L: this follows easily from the compactness. If $L(\alpha_n)$ tends to ∞ but $\alpha_n\in B_x$ then some subsequence converges to some α; but since L is linear hence continuous, we must then have $L(\alpha)=\infty$, which is absurd. Note that the maximum is achieved, since L is continuous and B_x is compact.

It is easy to see that x^* is a semi-norm; linearity on rays is almost immediate, and convexity simply follows from the fact that the max of a sum is less than the sum of the max's. It is also in fact a norm, since any non-zero X_i takes a non-zero value at some point α; some multiple of it has norm less then 1, and so it's absolute value is positive. It therefore has, as x, a compact unit ball B_{x^*}.

Our next task will be to show a much more interesting property of these norms:

Theorem: The unit balls B_x, B_{x^*} are convex polyhedra, i.e., they are each the intersection of a finite number of half-spaces.

Foliations and the Topology of 3-manifolds

Outline of class 20

Well, that was another awful lecture. Here's what I should have been saying.

Our goal for the day is to prove:

Proposition: The unit balls B_x, B_{x^*} are convex polyhedra, i.e., they are each the intersection of a finite number of half-spaces.

This is actually true for any norm on $\mathbf{R^n}$ which takes integer values on the integer lattice (think for example of the L^1 norm on $\mathbf{R^n}$). So in our proof we will just assume we have a norm $x:\mathbf{R^n}\to\mathbf{R}$ such that $x(\alpha)\in\mathbf{Z}^+$ for all $\alpha\in\mathbf{Z^n}$.

Lemma: If a,b$\in\mathbf{Z}^n$, then there is an N such that for all k$\geq$0, x((a + Nb) + kb)=x(a + Nb) + kx(b).

Proof: Let f(m) = x(a + (m+1)b) $-$ x(a +mb)$\in\mathbf{Z}^+$. By the convexity of x, f(m)$\leq$x(b); we claim that f is also an increasing function of m. This can be seen quite easily by writing

$$\begin{aligned} f(m+1)-f(m) &= (x(a + (m+2)b) - x(a + (m+1)b)) - (x(a + (m+1)b) - x(a + mb)) \\ &= (x(a+(m+2)b) + x(a + mb)) - 2(x(a + mb)) \\ &= (x(a+(m+2)b) + x(a + mb)) - x(2(a + mb)) \\ &= (x(a+(m+2)b) + x(a + mb)) - x((a + (m+2)b) + (a + mb)) \ , \end{aligned}$$

which is positive by the convexity of x.

Therefore, since f is increasing, integral, and bounded from above, it must eventually be constant; there is an N so that f(N+k)=C for all k$\geq$0, for some C. It only remains to see that C=x(b), since by induction x((a + Nb) + kb) $-$ x(a + Nb) = kC (which would give the result). But by dividing this equation by k we get

$$(*) \qquad x(\frac{a + Nb}{k} + b) - x(\frac{a + Nb}{k}) = C,$$

but by continuity of x, as k approaches ∞ ((a+Nb)/k approaches 0, so) the left-hand side approaches x(b)-0, so C=x(b).

Using this lemma we can see that x is affine linear on the line segment between b and b+(a+Nb) (see figure); this follows from the above calculation for k=1 and 2. For if, for some 0$\leq$t$\leq$1,

$$x(b+t(a+Nb))=x((1-t)b+t(b+(a+Nb)))< (1-t)x(b)+tx(b+(a+Nb))$$

(it must be less than or equal, by convexity), then

$$\begin{aligned} x(b+t(a+Nb))&< (1-t)x(b)+tx(b+(a+Nb)) \\ &\leq (1-t)x(b)+t(x(b)+x((a+Nb)))=x(b)+tx((a+Nb)), \qquad \text{so} \\ x(b+(a+Nb) &= x((b+t(a+Nb))+(1-t)(a+Nb)) \\ &\leq x(b+t(a+Nb))+(1-t)x(a+Nb) \\ &< (x(b)+tx(a+Nb))+(1-t)x(a+Nb) \\ &= x(b)+x(a+Nb) \end{aligned}$$

(by convexity and combining terms), contradicting (*) for k=1.

But then because x is linear on rays through the origin, it is easy to see that x is therefore linear on the cone on this line segment [b,a+(N+1)b] in $\mathbf{R}^n$. This is the basic building block of our construction. What we will do is show that, for any indivisible element b of $\mathbf{Z}^n$ (meaning it can be extended to a $\mathbf{Z}$- basis b=$a_1,a_2,\ldots,a_n$of $\mathbf{Z}^n$) x is linear on the convex span of some ($\mathbf{Z}$-)basis for $\mathbf{R}^n$ (meaning it is linear on the cone of the convex hull of the points in $\mathbf{R}^n$ that the basis vectors represent (or, equivalently, the set of linear combinations with non-negative coefficients)) which contains b (as a vertex). This means that x agrees with a linear function L_b on this span; but because x is integral on $\mathbf{Z}^n$, and hence on the basis vectors a_i, L_b is an <u>integral</u> linear function. This will be very relevant later on.

The idea is to do this by induction (on the number of basis vectors). We've just seen how to do the base case; starting with b and another basis vector a, by adding b to a enough times (i.e., replacing a with $a+(N+1)b$, which still gives a basis), x is linear on their convex span. If we have managed to show that x is linear on the convex span of $b=a_1,\ldots,a_k$, then by taking $b'=(1/k)(a_1+\cdots+a_k)$ (their barycenter), and applying the argument above to b' and a_{k+1}, we can find a <u>new</u> a_{k+1} (the old one with b' added to it enough times, so that the new a_{k+1} and the remaining a_i's still give a $\mathbf{Z}$-basis of $\mathbf{Z}^{\mathbf{k+1}}$) so that x is affine linear on the line segment from $b'$to a_{k+1}. But then a rather nice dirty trick allows us to see that x is then affine linear on the convex hull of $a_1, \ldots, a_{k+1}$ (and hence linear on their convex span). For if we pick any point $\alpha=\Sigma_{i=1}^{k+1}t_i a_i$ in the convex hull (so $0\leq t_i\leq 1$ for $1\leq i\leq k+1$ and their sum is 1), then

$$x(a_{k+1})\leq x(\Sigma_{i=1}^{k} t_i a_i)+x(t_{k+1} a_{k+1}) =\Sigma_{i=1}^{k+1} t_i x(a_i)$$

(by convexity and linearity of x on the convex span of the first k vectors). On the other hand, we can also write $\gamma=t\alpha+(1-t)\beta$, for some point γ on the line segment between $b'$and a_{k+1}, and some point β in the convex hull of $a_1,\ldots,a_k$ (this is easy to see if we just restrict our attention to the plane in $\mathbf{R}^\mathbf{n}$ containing b', α, and a_{k+1} - see figure (Oh - and thank you, Nathan)). Then from convexity it is easy to see that $x(\gamma)\leq tx(\alpha)+(1-t)x(\beta)$, so $x(\alpha)\geq 1/t(x(\gamma)-(1-t)x(\beta)$. But since γ and β are both in places where we know x is linear, we can write them as a sum of multiples of $x(a_i)$, and so doing this and combining terms (since the resulting coefficients must be the same as the coefficients of α as a sum of the a_i(i.e., $\alpha=(1/t)\gamma-((1-t)/t))\beta$), we get that $x(\alpha)\geq \Sigma_{i=1}^{k+1}t_i x(a_i)$, giving equality, and hence affine linearity, on the convex hull of $a_1,\ldots a_{k+1}$.

Therefore, by induction, we have shown that every indivisible b can be extended to a basis of $\mathbf{Z}^\mathbf{n}$ (and hence of $\mathbf{R}^\mathbf{n}$) so that x agrees with a linear function L_b on their convex span $C=C_b$.

Notice that we also get that x agrees with $-L_b$ on the convex span of $-a_1,\ldots,-a_n$; both functions change sign. But now since x and L_b agree on this cone, they are both equal to 1 at the same points of this cone. For L_b, these points consist of the intersection of an affine hyperplane P_b with the convex cone C_b. The hyperplane determines a half space $H_b=\{\alpha\in\mathbf{R}^\mathbf{n} : L_b \leq 1\}$.

Lemma: $B_x\subseteq H_b$, i.e., if $x(\alpha)\leq 1$, then $L_b(\alpha)\leq 1$.

Proof: Suppose not; suppose that for some $\alpha\in\mathbf{R}^\mathbf{n}$, $x(\alpha)\leq 1$, but $L_b(\alpha)>1$. Let $\beta=(1/n)(a_1+\cdots a_n)$ be the barycenter of our affine simplex, and let $\gamma=\beta/x(\beta)$, so $x(\gamma)=1$. Then travelling along the line segment from γ to α (see figure), x of every point is ≤ 1 by convexity (since this is true at the endpoints), but since $L_b(\gamma)=1$ (since x and L_b agree in the cone over the simplex) and $L_b(\alpha)>1$ by hypothesis, every point on the line segment has $L_b>1$, since L_b is linear. But this is absurd, since γ is in the interior of the cone where x and L_b agree (but they don't agree!).

Note that this also says that if $x(-\alpha)=x(\alpha)\leq 1$, then $L_b(-\alpha)=-L_b(\alpha)\leq 1$, i.e., $x(\alpha)\leq 1$ implies $|L_b(\alpha)|\leq 1$. Consequently, if we think of L_b as an element of the dual space $\mathrm{Hom}(H_2(M,\partial M;\mathbf{R}),\mathbf{R})$, it has $x^*(L_b)\leq 1$. In fact, since L_b agrees with x on lots of rays (namely any ray in the convex cone) on which some point of x-norm 1 lives, $L_b(\alpha)=1$ for some point with $x(\alpha)=1$, so $x^*(L_b)=1$. So all of the L_b's live in the unit ball B_{x^*} of the dual norm, which is compact. But since each of the L_b are integral linear, and hence correspond to the integer lattice in the dual vector space (we assume we chose the dual basis for the vector space), as b ranges over all indivisible elements of $\mathbf{R^n}$ L_b actually ranges over only finitely-many linear functions $L_{b1},\ldots,L_{bk}$.

This allows us to easily see that $B_x=H_{b1}\cap\ldots\cap H_{bk}$. B_x is contained in the intersection by the lemma and induction. To see the other containment, suppose we had $\alpha\in H_{b1}\cap\ldots\cap H_{bk}$ (so $L_{bi}(\alpha\leq 1$ for all i). Choose a sequence $(\alpha_n)_{n=1}^{\infty}$ in $\mathbf{Q^n}$ converging to α. Each α_n has a multiple (take the least common denomenator of their denomenators!) which is an indivisible point of its integer lattice, so there is a corresponding linear function in the finite collection above. Choose a subsequence (which we still denote $(\alpha_n)_{n=1}^{\infty}$) so that the corresponding linear functions L_β are all the same. Then $x(\alpha_n)=L_\beta(\alpha_n)$ for all n; but since $L_\beta(\alpha_n)$ converges to $L_\beta(\alpha)\leq 1$ (by hypothesis), we conclude that (since $L_\beta(\alpha_n)=x(\alpha_n)$ also converges to $x(\alpha)$, $x(\alpha)=L_\beta(\alpha)$, so) $x(\alpha)\leq 1$, as desired.

For B_{x^*}, we have seen that the linear functions L_{bi}, i=1,…,k are elements of the dual space of x^*-norm 1, so since B_{x^*} is convex (by the convexity of the dual norm), B_{x^*} contains the convex hull of these linear functions. But we claim that in fact the opposite is true; their convex hull equals B_{x^*} (this is known as the dual of the polyhedron B_x). To see this, pick a vertex v of the polyhedron B_x and look at the linear functions (call them $L_{b1},\ldots,L_{bk}$, for lack of a better name) corresponding to the (affine) hyperplanes that contain the (top-dimensional) faces that contain v. Because v is a vertex, the intersection of these hyperplanes is just $\{v\}$, which implies that the linear functions $L_{b1},\ldots,L_{bk}$ span the dual space. These functions all have $L_i(v)=1$ (since it lies in each hyperplane), and therefore any convex linear combination $L=\Sigma_{i=1}^{k}t_iL_i$ ($0\leq t_i\leq 1$, $\Sigma t_i=1$) has $L(v)=\Sigma_{i=1}^{k}t_iL_{bi}(v)=\Sigma t_i=1$, while $x^*(L)\leq\Sigma_{i=1}^{k}t_ix(L_{bi})=\Sigma t_i=1$ (by convexity), so (since $x(v)=1$) $x^*(L)=1$. In words, the convex hull of the L_{bi} form an (n-1)-dimensional (since the L_{bi} span) face of the boundary of B_{x^*} but since any point outside of the convex hull of all of the L_{bi}'s would lie over one of these faces (this is because the convex hull contains 0 (being the dual of a polyhedron), and any face of B_{x^*} corresponds to a vertex of B_x (the dual of the dual ...)), and hence (by linearity on rays) would have x^*-norm larger than 1, we have that B_{x^*} is contained in the convex hull.

That finishes our proof. Next time we will compute some specific examples for M = some link complements, and then start to explore how to use this machinery to construct foliations with norm-minimizing leaves.

Foliations and the Topology of 3-manifolds

Outline of class 21

We've now seen that if x is a norm, then B_x and B_{x^*} are convex polyhedra, and B_{x^*} is in fact the convex hull of the linear functions defining the faces of B_x. This is actually true in general, for x a semi-norm; this can be seen fairly quickly by looking at the norm induced from x. The set of points where x is 0 is a linear subspace X of $\mathbf{R^n}$, so decends to a norm on the quotient vector space $\mathbf{R^n}/X$. This is still an integral norm (properly interpreted) - if X has dimension k, then X intersects $\mathbf{Z^n}$ in a lattice of rank k (this follows easily from the fact that X is the $\mathbf{R}$-span of the points in the integer lattice on which x is zero). Therefore $\mathbf{R^n}/X$ still contains a naturally-defined integer lattice (the image of the one in $\mathbf{R^n}$) on which x is integral. So the proof we gave goes through to show that the unit ball downstairs is a convex polyhedron, so the unit ball upstairs is its inverse image under a linear map, hence a convex (though non-compact) polyhedron.

Exercise: Show that the vertices of B_x each have rational coordinates (Hint: each is determined by what top-dimensional faces it sits in, which gives linear constraints).

Next we will compute some specific examples of x and B_x for some knots and links. In general this can be a nearly impossible task, since we would need to be able to determine that least genus surface representing homology classes, which in practice is a very difficult task. But we will see that if our manifold has many symmetries, we can exploit this to help us in determining the norms of various elements, since they are reflected in symmetries of the norm itself (and hence of its unit ball).

We will draw our examples from link exteriors in S^3. (Consequently, we will take as known some of the basic machinery of knot theory). Using some duality theory from homology, we have that, for any link $L=L_1\cup\cdots\cup L_k$ in S^3, $H_2(S^3\setminus N(L),\partial N(L))\cong H^1(S^3\setminus N(L))\cong H_1(N(L))$, where this last isomorphisnm is by Alexander Duality. In fact, the composite isomorphism is rather easy to describe - a 2-dimensional homology class is sent to the homology class in N(L) of its boundary. But $H_1N(L)\cong\mathbf{Z^k}$, on generator for each component of the link, and if we orient each link, we can make a natural identification of $H_1N(L)$ with $\mathbf{Z^k}$. by choosing as generator the oriented circle itself.

Now for knots, the Thurston norm is fairly boring. It is determined by the value of x(K) (abusing notation and identifying a 2-dimensional homology class with its boundary), so is a norm iff $x(K)\neq 0$. But the only surface with $\chi_-(S)=0$ and homologically having boundary K is a disk (annuli have 2 boundary components (parallel in $\partial N(K)$), so could only represent an even multiple of the generator), so $x(K)=0$ iff K is the unknot. But then if K is not the unknot, then $\pm K/x(K)$ are the only elements of norm 1, so $B_x=[-K/x(K),K/x(K)]\subseteq\mathbf{R}$.

So if we really want to get some interesting examples, we need a link with more than one component. But if we also want to actually be able to calculate x effectively on $H_1(L)$, in most cases we are going to need a link with alot of symmetry.

To see how symmetry might help, suppose $h:M\to M$ is a self-homeomorphism of a manifold M. Then because of the way x is defined, $x(\alpha)=x(h_*(\alpha))$ for all $\alpha\in H_2(M,\partial M;\mathbf{R})$. This is because if S is a surface of minimal χ_- representing α, then h(S) (has the same χ_- and) is a minimal χ_- representative for $h_*(\alpha)$; otherwise, $h_*(\alpha)$ has a surface T of smaller χ_-, and then $h^{-1}(T)$ represents α amd has smaller χ_- than S, a contradiction. Therefore, the symmetry of the manifold is reflected in a symmetry of the norm, and hence in a symmetry of the unit ball. This fact can actually be useful for computational purposes.

Take for the example the Whitehead link L (see figure below), with an orientation on each component. Label the components L_1,L_2 respectively. Then there are several symmetries of L that we can fairly easily find. By rotating 180 degrees around a line through the center of the figure and our eye, we can take L_1 to L_1 and L_2 to $-L_2$ (meaning L_2 with the opposite orientation). By rotating about the vertical line through the center, we can take (L_1,L_2) to $(-L_1,L_2)$. Finally, by turning our picture of L inside out (turning L_2 into the round circle) we can see that we can take (L_1,L_2) to (L_2,L_1). Therefore the group generated by these symmetries (which is the dihedral group D_4) acts as symmetries of x. So we can permute L_1 and L_2 and introduce signs at will, without affecting the value of x. This allows us to effectively calculate x.

Start with $x(L_1)$ $(=x(L_2))$. We need to find a surface whose boundary represents L_1 in $H_1(N(L))$. Inspection quickly turns up one of the surfaces S,T below; the first is a once-punctured torus, and the second is a twice-punctured disk. The second is really no better than the first - it's obtained by pinching off the annular piece following the knot. But the second can be a very useful construction to remember in other situations - it represents the same homology class, because the two small loops are meridians, so bound disks in N(L), so are zero in $H_1(N(L))$ - they contribute nothing. So we can introduce them (or live with them) when looking for our surfaces without introducing any problems.

$\chi_-(S)=\chi_-(T)=1$, so we immediately know that $x(L_1)\leq 1$. But if $x(L_1)=0$, then there is a surface F representing L_1 with $\chi_-(F)=0$, hence is a disk or an annulus. But L_1 can't bound a disk, since this would imply the L_1 could be shrunk to a point in the complement of L_2, and hence L is the unlink (which we take as given (i.e., use other machinery) to be untrue). It also can't be an annulus - if it were, one component is L_1, so the other must be a meridian of $N(L_2)$ (a trivial loop in $\partial N(L)$ could be pushed off, reducing to the case of a disk). If we glue on a meridian disk, we then get a surface (a disk) in S^3 bounded by L_1 and intersecting L_2 exactly once. This, however, implies that L_1 and L_2 have linking nmuber ± 1; but the original surface S (or T with the two meridians capped off by disks) shows that they have linking number 0, a contradiction. So $x(L_1)=x(L_2)=1$.

From this we can conclude that $x(L_1+L_2)=2$; for if not, convexity says it must be 1 or 0. 1, however, can be quickly ruled out, because any surface representing L_1+L_2 must have an even number of (essential) boundary components (hence an even χ_-). This is because on each boundary torus we have a collection of parallel loops, whose sum represents a generator; an even number would represent an even multiple of the generator. So there are an odd number of loops on each component, hence an even total number of them. So we would have to have $x(L_1+L_2)=0$; but since we already know that $x(L_1+L_2)=x(L_1-L_2)$, we would then have that $2=x(2L_1)\leq x(L_1+L_2)+x(L_1-L_2)=0+0=0$, a contradiction.

This tells us then that $x((1/2)\pm L_1\pm L_2)=1$ (where the two signs may be chosen independently), so we have found eight points on the boundary of the unit ball B_x - see the figure below. But it is easy to see that the only convex set in R^2 that contains these points on its boundary is the diamond - if it contained any other point (by symmetry, in the first quadrant, say), then by taking the convex hull of the diamond and this point, the point $(1/2,1/2)$ would end up in the interior of B_x, a contradiction.

With this information, we can easily calculate the Thurston norm of any element of $H_2(M,\partial M;\mathbf{R})$; $x(aL_1+bL_2)=|a|+|b|$ (prove it!).

Notice that we in fact proved that B_x as a polyhedron, in this case, instead of using it to help us find out what it looked like. This is something that can be used in general. If you already have (somehow) a collection of points which you know are in ∂B_x, and they include both a collection (the vertices) of points P and a point in the interior of each of the top-dimensional faces of the convex hull of P, then since the only convex set which contains all of these points in its boundary IS the convex hull of P, this must be exactly the unit ball of the norm.

Let's apply this same sort of reasoning to another example, the Borromean rings (see figure). Each component bounds a once punctured torus that misses the others (which we could also think of as twice-punctured disks) - see figure - so since the link is (we assume) non-trivial, and he linking numbers of each pair of components is 0, we can conclude as in our previous example that $x(L_1)=x(L_2)=x(L_3)=1$.

Again, the Borromean rings has a great deal of symmetry, which we can exploit in our calculations. If we rotate the picture around its center, we can permute the components $(1,2,3)\mapsto(2,3,1)$ and to $(3,1,2)$. By rotating around lines in the plane which are vertical, and inclined 120 degrees, we get the permutations $(1,2,3)\mapsto(-1,-3,-2)$, $(-3,-2,-1)$, and $(-2,-1,-3)$. Finally, if we perform an inversion (in R^3) the link around the center of one of the components (say L_1), and then reflect in the plane of the link projection and rotate it back onto itself, we find we have achieved the permutation (1,2,3)$\mapsto$(1,-2,-3). Composing all of them we can then get (1,2,3)$\mapsto$(-1,-2,-3) , and <u>then</u> composing with another one above gets us to (1,3,2). So we can achieve any rotation or transposition or change of sign, hence our norm will have the corresponding symmetries.

Now $x(L_1+L_2)=2$; this is similar to the calculation we gave before. It can't be zero since then symmetry implies $2x(L_1)\leq x(L_1+lt)+x(L_1-L_2)=0$. Also, it can't be 1; since a χ_--minimzing surface S has at least two ∂-components (and an even number of components altogether on $\partial N(L_1)\cap\partial N(L_2)$), and an odd number of ∂-components (since $\chi(S)$ is odd), there must be a meridian on L_3; by capping off with a disk, we get a surface bounded by $L_1\cup L_2$ intersecting L_3 exactly once, again contradicting linking number 0.

Finally, using this and symmetry, we can conclude that $x(L)=x(L_1+L_2+L_3)=3$, for otherwise convexity (and the fact that, as before, a representative has to have an odd number of essential boundary components (an odd number around each component) $x(L)=1$; but then using symmetry, $4=x(2(L_1+L_2))\leq x(L_1+L_2+L_3)+x(L_1+L_2-L_3)=1+1=2$, a contradiction.

From this we can deduce the shape of B_x; each of the $\pm L_i$ lie on ∂B_x, and their convex hull is the octahedron. Burt the points $(1/3)(\pm L_1\pm L_2\pm L_3)$ (signs are independent) each lie in the interior of a different face of the octahedron, and all lie in ∂B_x; so our argument before tells us that the octahedron$=B_x$. From this we can quickly determine x itself; $x(aL_1+bL_2+cL_3)=|a|+|b|+|c|$. Also, since the faces of the octahedron correspond to the vertices of B_{x^*} (the vertex isd the linear map equal to one on the face), we can determine the unit ball B_{x^*} as well. The linear map is the inner product with the normal to the plane containing the face (suitably scaled); but these normals are easily seen to be $(1/3)(\pm1,\pm1,\pm1,)$ (signs independent), and the scaling factor (since the norm of this vector is $1/\sqrt{3}$) is $\sqrt{3}$, so it easily follows that $B_{x^*}=[-\sqrt{3},\sqrt{3}]^3$.

Next time we will do a somewhat more complicated example (another 3-component link), and then start looking at the mechanics behind Gabai's construction of foliations using the Thurston norm.

Foliations and the Topology of 3-manifolds

Outline of class 22

We'll do one more example of computing the Thurston norm, before moving on to more theoretical business.

Let L be the three-link chain pictured below. As with our previous examples, we will exploit some symmetries of this link to help us compute the norm: with the orientations we have given it, rotating about the center of the picture affects the permutations (1,2,3)$\mapsto$(2,3,1) and (3,1,2). Reflecting it in the lines in the plane and through the center give the permutations $(1,2,3) \mapsto (-1,-3,-2)$, $(-3,-2,-1)$, and $(-2,-1,-3)$. By drawing the link slightly differently (see below) and rotating around the center, we get the permutation $(1,2,3) \mapsto (1,3,2)$. Finally, composing this with on of the transpositions above, we get $(-1,-2,-3)$. From these it is easy to see that any permutation that changes no sign or changes all of the signs can be achieved as a symmetry of the link, and therefore reflects a symmetry of the Thurston norm x.

Now let's calculate. As before, $x(L_1)=1$; a surface with this χ_- is given below, so if $x(L_1) \neq 1$, then $x(L_1)=0$, so L_1 could be represented by a disk or annulus. But, as before, a disk contradicts the non-triviality of the link (or, if you prefer, the fact that the linking number of pairs of components are all 1, in this case), while an annulus would give a disk spanning L_1 and hitting only one of the other components, contradicting linking number 1 with the other component.

In this case, however, we now jump out of order and calculate $x(L_1+L_2+L_3)$ next. A surface with $\chi_-=1$ is given below, and since the Euler characteristic of a spanning surface for L must be odd (there are an odd number of boundary components around each component) and negative (if $=1$, some component would be a disk, contradicting non-triviality again), $x(L)$ must be at least 1, so $x(L)=1$.

Now we can show $x(L_1+L_2)=2$; if not, it is 1 or 0. If 1, then the χ_--minimizing surface would have to be a twice-punctured disk (since it has to have at least 2 components); but then one component (around L_3) is a meridian, so the surface gives an annulus hitting L_3 once, contradicting that $L_1 \cup L_2$ has linking number 2 with L_3. On the other hand, if $x(L_1+L_2)=0$, then exploiting the symmetries we have of the link, we would have
$2=2x(L_1+L_2+L_3)=x((L_1+L_2)+(L_2+L_3)+(L_3+L_1)) \leq x(L_1+L_2)+x(L_2+L_3)+x(L_3+L_1)=0$, a contradiction. So $x(L_1+L_2)=2$. But from this it then follows that $x(L_1+L_2-L_3)=3$; for otherwise (since a representing surface must again have an odd number of boundary components) it equals 1, and then

$$4=2x(L_1+L_2)=x((L_1+L_2)+(L_1+L_2)) \leq x(L_1+L_2+L_3)+x(L_1+L_2-L_3)=1+1=2,$$

a contradiction. Therefore x(L_1+ł$-L_3$)=3, allowing us to finally prove that $x(L_1-L_2)=2$, because (using the symmetries)

$$3=x(L_1-L_2+L_3) \leq x(L_1-L_2)+x(L_3)=x(L_1-L_2)+1,$$

so $x(L_1-L_2) \geq 2$, while convexity says it is ≤ 2.

Therefore $x(\pm L_1 \pm L_2 \pm L_3)=1$ if all of the signs are the same, otherwise it is $=3$; and $x(L_i \pm L_j)=2$ for all $i \neq j$. If we plot all of the points this then guarantees have norm 1 (see below), we find that their convex hull is basically the octagon that formed the unit ball for the Borromean rings (which is what we would have gotten here if $x(L_1+L_2+L_3)=3$) together with two tetrahedra stuck onto it in the all-positive and all-negative octants. The amazing thing, however, is that is gives us a convex polyhedron in which we have already found points of norm 1 in the interior of each face! The point is that, for example, the triangle with vertices (0,-1,0), (1,0,0), and (0,0,1) and the triangle with vertices (1,0,0), (0,0,1), and (1,1,1) together form a planar rhombus (by checking that the lines between the pairs ((1,0,0), (0,0,1)) and ((0,-1,0), (1,1,1)) intersect (at (1/2,0,1/2)). The resulting polyhedron has six faces (six faces of the octahedron pair off with the faces of the tetrahedra, and two faces were lost when the tetrahedra were glued on). We have also found points of norm 1 (the intersection of the two diagonals are of the form $(1/2)(L_i \pm L_j)$) in the interior of each face! So this polyhedron is the unit ball! It's basically a 'rhombic cube' - it's what you would get if you took two opposite corners of a cube and pulled them until the short diagonal of a face equalled the side length.

So that's one final, somewhat more interesting example.

What we will start to do now is to begin an attack on Gabai's converse to Thurston's theorem:

Theorem (Gabai): If M is a compact, orientable 3-manifold, with ∂M = tori, and if $(S,\partial S)\subseteq(M,\partial M)$ is a surface which is norm minimizing in its homology class, then there exists a taut foliation $\mathcal{F}$ of M, transverse to ∂M, such that S is a leaf of $\mathcal{F}$.

One very important application of this theorem is to computing the genera of knots and links in S^3. If L is a link in S^3, and $M = S^3 \setminus \mathrm{int}\ N(L)$, then there is a surface (representing L under Alexander duality) S with ∂S = a loop in each ∂-component of M (we often think of it as a surface in S^3 with boundary equal to L), known as a spanning surface for L. One interesting problem is to compute, for a given link L, the least possible genus of such a surface for L, and construct a representative. In the language of Thurston norms, we wish to calculate x(L), and find a norm-minimizing representative. Gabai's theorem gives an often effective way to do this; if you have a likely candidate, then by foliating the complement we can invoke Thurston's theorem to prove it is norm-minimizing; Gabai's theorem says that if we are right, then we can (in principle, and often in practice) construct such a foliation.

To understand Gabai's proof (we will not go through it in all of its generality, but will do so in a special, but very useful, case), we need to introduce some new concepts.

Definition: A **sutured manifold** is a pair (M,γ), where M is a compact orientable 3-manifold, and $\gamma\subseteq\partial M$ is a collection of disjoint annuli and tori, denoted $A(\gamma)$ and $T(\gamma)$. Each annulus A comes with an orientation, in the form of an orientation to its core circle, called the suture, and denoted s(A); the collection of sutures is denoted $s(\gamma)$. Each component of the rest of the boundary, $R(\gamma)=\partial M\setminus\mathrm{int}(\gamma)$, can be given a (normal) orientation, allowing us to split $R(\gamma)$ into two pieces; $R_+(\gamma)$ consist of those pieces for which the normal points out of M, and $R_-(\gamma)$ consists of the pieces with inward-pointing normal. These normals, however, must be chosen in a way consistent with the orientations of the sutures - if we orient a component of $R_\pm(\gamma)$ (using the normal orientation as third vector of an orientation for M), this induces an orientation on its boundary components; these must agree with the orientations on the suture in the bordering annulus. See the figure below.

This rather long-winded definition is obscuring what is really a rather simple idea; the motivating picture is what you get when you split a link exterior open along an orientable spanning surface S (shown below). The surface S becomes two surfaces in the boundary of the resulting 3-manifold M=$S^3\setminus$int N(L) split open along S, with annuli running between the pairs of boundary components (coming from ∂N(L); this is γ). The orientation on S induces a normal orientation, which gives normal orientations for the two copies of S in ∂M. Finally, if S homologically represents the link L, then the orientation it induces on it's boundary is the same as the orientation on the link; this orientation on the link gives the orientations of the sutures which are compatible with the normal orientations on $S_\pm = R_\pm(\gamma)$.

One thing to notice: the orientation conventions of a sutured manifold require that as you cross a suture you must pass from R_+ to R_- (or vice-versa) - the orientation on the boundary reverses itself as you cross a suture (from the point of view of ∂M).

Now it is a fact that, for any compact, orientable 3-manifold M and any codimension-0 submanifold N of ∂M, we can define a Thurston norm on H_2(M,N), just as we did for H_2(M,∂M). One just needs to know that any homology class in this relative group can , like before, be represented as a compact, embedded surface S with $\partial S \subseteq N$. I've never actually seen a proof of this, but since I'm only mentioning it mostly for motivational purposes anyway, this small lack will not bother us much. Then we came make the definition:

Definition: A sutured manifold (M,γ) is called **taut** if $R_+(\gamma)$ and $R_-(\gamma)$ are both norm-minimizing in their homology classes.

Here the motivating example is that the sutured manifold described above is taut, if S is norm minimizing in its homology class in $H_2(S^3\setminus$int N(L),∂N(L)) (proof next time). Then the theorem Gabai actually proves is:

Theorem (Gabai): A sutured manifold (M,γ) is taut iff there exists a taut foliation $\mathcal{F}$ of M, transverse to γ, with R(γ) = union of leaves of $\mathcal{F}$.

Next time we will discuss some more of the basic ideas that go into the proof; it is basically an induction, using the notions of a decomposing surface and a sutured manifold hierarchy.

Foliations and the Topology of 3-manifolds

Outline of class 23

Last time we talked about sutured manifolds and taut sutured manifolds. Today we will describe what they are good for (and how one might find one!). Since from this point forward much of what we do will rely heavily on pictures, I'm going to start sticking them at the end.

Gabai's construction employs induction, using the notion of a decomposing surface:

Definition: A **decomposing surface** S for a sutured manifold (M,γ) is a ((normally) oriented, properly embedded) surface in M such that

(1) For every component T of T(γ), S$\cap$T consists of parallel, coherently-oriented circles;

(2) For every component A of A(γ), S$\cap$A consists either of parallel circles coherently-oriented with one another and the suture in A, or arcs running straight across the annulus A; see Figure 1.

Figure 1

The main thing one does with a decomposing surface is to decompose along it; we can create a new manifold M$'$ by splitting M open along S. The point, though, is that because of the conventions on how the boundary of S meets the sutures, we can in fact get a new sutured manifold, as follows:

In case (1) and the first half of case (2) above, S splits T and A into annuli (see Figure 2). Using the orientations on the boundary of S, we can assign an orientation to the cores of these annuli (in case (2), this orientation agrees with the one already present). Also (locally, at least) we can use the normal orientation to assign a normal orientation to all of the new boundary components, so that the new annuli are in fact sutures for M$'$. In the second half of case (2) we can do something similar, it is just a little bit more subtle. The picture before and after splitting is as in Figure 3; what we do is use the normal orientations on each of the pieces to decide where the sutures should be. The sutures are supposed to separate the parts of the boundary where the normal orientations are reversed, and where we look at the picture, we see that this should include the two pieces we originally started with, as well as two pieces (one top, one bottom) where s met R_+ and R_-. Notice that if we had given S the opposite orientation (assuming we could, i.e., S didn't also meet an annular suture in parallel circles (which forces an orientation upon us, to get coherency) our choices here would be the other two arcs, instead. These four pieces fit together to give us two pieces of our new sutures γ'. We write (M,γ)$\overset{D}{\Rightarrow}$(M$'$,γ').

Figure 2 Figure 3

Gabai's proof that taut sutured manifolds admit taut foliations then breaks into two main pieces. First he shows that a taut sutured manifold has a (taut) **sutured manifold decomposition**:

Definition: A (taut) sutured manifold decomposition of a sutured manifold (M,γ) is a sequence of decomposing surfaces

$$(M,\gamma)\overset{S_1}{\Rightarrow}(M_1,\gamma_1)\overset{S_2}{\Rightarrow}\dots \quad \dots\overset{S_n}{\Rightarrow}(M_n,\gamma_n)$$

where each sutured manifold is taut, and $(M_n, \gamma_n) = (\coprod D^2 \times I, \coprod \partial D^2 \times I)$. The sequence of sutured manifolds is called a sutured manifold hierarchy for M.

Then Gabai shows that one can use the sutured manifold hierarchy to tautly foliate M. One works back up the hierarchy from the bottom; it is easy to see how to tautly foliate (M_n, γ_n) - foliate each 3-ball by parallel disks. This is transverse to the sutures, and contains the surfaces R_+, R_- as leaves. Then by using the picture of how (M_k, γ_k) is obtained from (M_{k+1}, γ_{k+1}) (by gluing back together), one can extend the foliation on (M_{k+1}, γ_{k+1}) to one on (M_k, γ_k) (with the same properties) by a process of spinning, shown diagrammatically in Figure 4. One basically take s an infinite number of parallel copies of the R_+, R_- in M_{k+1}, limiting on the real ones; then by cutting them open along S_{k+1} and shifting by one, we can basically turm M_k into M_{k+1} with an infinite product glued on. Gabai shows that by making careful choices, one can extend the foliation on M_{k+1} to one on M_k, by gluing on a foliation of this infinite product. This is a very delicate process, which we will only prove in a special (although highly useful) case, that of a **disk decomposition**.

Figure 4

Definition: A disk decomposition for a sutured manifold is a sutured manifold in which all of the decomposing surfaces are disks. (We do not require that all of the sutured manifolds in the resulting hierarchy be taut; this is actually implied by the existence of the decomposition.) A sutured manifold which has a disk decomposition is called **completely disk decomposeable**.

As an example, take the (Hopf) link L pictured in Figure 5 (which we think of as sitting in S^3). We can easily find a spanning surface for L, and if we imagine thickening up this (annular) surface, we get a solid torus (whose complement, which is actually our sutured manifold, is also a solid torus). We can then draw the (cores of the) resulting annular sutures on this solid torus, to gives us an inside-out view of our sutured manifold. Remembering that the disks we want to find should lie outside of the solid torus we see in from of us, it's not hard to find a candidate for a decomposing disk. All that remains is to figure out what it looks like when we decompose along this disk. It's clear that the boundary of our new sutured manifold will be a sphere (and our sutured manifold will be a solid torus cut open along a meridian disk, i.e., a ball) - we must then figure out what the new sutures look like. But this is actually not to hard to figure out. The boundary of the (oriented) disk comes with its own orientation, as do the two original sutures, and the new sutures inherit these same orientations, so there is really only one way to glue the pieces together in a way that respects orientations - see the figure. Inspection shows that the resulting suture is connected, so the sutured manifold we get is actually $D^2 \times I$. Therefore the sutured manifold we started with is disk decomposable.

Figure 5

The general picture we get when decomposing along a disk is not much worse than what we drew above - see Figure 6. The picture is entirely local; basically the old sutures get surgered along the disk D, according to the orientation of the disk. By choosing the opposite orientation, we get what is in essence a dual picture. These two different decompositions can create drastic differences in the global behavior of the sutures. In our constructions it will often be the case that one choice gives us a much more well-behaved collection of sutures.

Figure 6

What we will show is that if we have a disk decomposition for a sutured manifold (M,γ), then we can inductively build up a taut foliation (as in Gabai's general argument) . Therefore, if we started by splitting the exterior of a knot or link along a spanning surface, we will have shown that the spanning surface is norm minimizing in its homology class (which is almost - but not quite - the same as saying it is a least-genus spanning surface for the link). But just as important, we will also produce a family of examples of spanning surfaces for several classes of knots and links which we can (relatively) easily show are disk decomposable, thereby allowing us, for example, to compute the genera of these knots and links. We will in fact do this first (in order to give us some reason for actually going through with the construction of the foliations!); the construction of the spanning surfaces themselves will be by Seifert's algorithm.

Given an (oriented) diagram for a knot or link L , Seifert described a process for building an embedded surface in S^3 whose boundary is L. We start by projecting L onto the plane, creating an (oriented) 4-valent graph, like in Figure 7. Using the orientations on the edges, we break the graph at the vertices (by matching them in adjacent pairs head-to-tail) to create a collection of disjoint oriented loops. We imagine bounding each loop by an embedded (oriented) disk, by pushing the loops into different levels, as necessary. Then we glue rectangles to these disks, in the form of half-twisted bands (which is our way of re-introducing the vertices of the link projection) to recover a surface whose boundary is L.

Figure 7

There are two things that we can notice about this construction. The first is that the surface that we build this way is always orientable. (N ormal) orientations on each of the disks can be induced from the orientations of their boundaries, as can the half-twisted bands (two opposite edges of the rectangle are coherently oriented (the ones glued to the disks), which determine the orientation). These normal orientations agree at the point where any one is glued to another (more or less by definition!), and hence determine a global (normal) orientation.

The other is that the genus of the resulting surface really doesn't matter so much on the link we started with, as it does the oriented graph that it creates. For if we were to change a crossing in our link projection, the oriented graph remains the same, and so the disjoints loop (and the disks they bound) do, too. The only thing that changes is the direction of twisting given to the rectangle that we glue on a give the spanning surface. But the direction of twisting can be changed by cutting along the middle of the rectangle and giving a full twist; consequently, theF resulting surfaces are homeomorphic, and hence have the same genus.

This last fact allows us to easily see that Seifert's algorithm will not always build a minimal genus spanning surface for the link (or, more precisely, a norm-minimizing surface for the homology class it represents). For example, if we take our picture of the figure-8 knot above in Figure 7, and change one crossing, we get a projection of the unknot. But it is easy to see that the algorithm run on the figure-8 knot produces a surface (of Euler characteristic -1, hence) of genus 1, is it also does for the unknot. But the genus of the unknot is 0 - it can be spanned by a disk.

Figure 8 (no joke intended)

However, we shall see next time that, in many cases, for the right projection of the link Seifert's algorithm does build a norm-minimizing surface spanning the link. Such projections include all alternating projections, as well as the standard pictures of pretzel knots (with a few exceptions). We will show this by showing that the resulting surface is disk decomposable.

Foliations and the Topology of 3-manifolds

Outline of class 24

Well, today's proofs will consist almost entirely of pictures.

Our goal is to prove:

Theorrem: If L is a (non-split) alternating link projection, then the surface we get by running Seifert's algorithm on L is disk decomposable.

The proof naturally breaks down into two cases. The first will be dealt with today - we'll do the other case next time. In each case the idea of the proof is to argue by induction on the Euler characteristic of the surface created - if χ(F)=1, then F is a disk, and $S^3\setminus$int N(F)=$D^2\times$I, which is clearly disk decomposable (the non-splitness assumption is just to insure that the surface we build using the algorithm is connected).

Basically, if we run Seifert's algorithm and look at the disjoint circles that are created halfway through (see Figure 1), one of two things happens. Either the circles can be filled in with disks in the plane (well, actually, 2-sphere) of the projection, which do not intersect one another (the 'Seifert circles' are not nested), or they cannot be (the Seifert circles are nested). We will deal with the first case first.

Figure 1

Now, whether or not the Seifert circles are nested really doesn't depend on the link, but just the projection of it onto the plane (i.e., on the underlying oriented graph in the plane). We use the alternating hypothesis, however, to tell us that the half-twisted bands that are attached to the Seifert disks all twist in the same direction - they either all have a right-handed twist or all have a left-handed twist (Figure 2). This is because as we travel around the boundary of a Seifert disk, we always find bands with the same twist (otherwise, consecutive bands of opposite twist would give consecutive overcrossings, say).

Now associated to our Seifert surface there is (in this case) a graph Γ (which we can think of as being both in the surface F and in the projection plane) - we put a vertex in the center of each Seifert disk, and an edge across each half-twisted band (see Figure 3). By altering our link projection if necessary, we can also arrange that

(1) The graph Γ has no hanging edges (Figure 4), and more generally

(2) No edge of Γ separates Γ (Figure 5).

In the first case, a hanging edge means a Seifert disk with only on band attached to it; by flipping the disk over, we can amalgamate two disks and a band into a single disk. Clearly, the projection remains alternating. In the second case, there is a loop in the plane which hits Γ once in that edge. this loop separates the plane into two disks, and by flipping the entire link projection in one of them over, we again amalgamate two disks and an edge (the offending one) into a single disk. Again it is not hard to see the resulting projection of our link is still alternating, and further that the resulting surface is the same as we would get by running Seifert's algorithm again on the new projection (Figure 6). Since these actions in both cases reduce the number of crossings in our projections, we eventually can't do either. (Those of you attending Cameron's class no doubt recognize this as the process of getting a reduced diagram for an alternating link.)

Figures 3 through 6

But now if we look at our graph Γ, it turns out that any of the disks that it cuts the sphere into will serve as an appropriate disk to start our disk decomposition on. (We went through the trouble of reducing the diagram first just so that the boundary of each of these disks do not cross the same edge of the graph twice - otherwise an arc joining those two 'faces' of the disk gives a loop appropriate for (2)).

In terms of our Seifert surface, this disk looks like Figure 7. Remember that alternating implies that the bands all twist in the same direction - our disk looks like a 2n-gon (since F is oriented - the orientation reverses as you cross each band, so must change an even number of times). Thickening up the picture, we get Figure 8. Now if we choose our favorite orientation (either one will do), and perform the decomposition, locally we get a picture like Figure 9. It's just a matter of breaking our old sutures where they hit the disk, and then following in the direction of the boundary of the disk to reglue them. But if we pull things a little tight, we then get a picture like Figure 10, and it is easy to see that this picture comes from a thickened-up surface (locally) - see Figure 11. The point is that this decomposition really amounts to removing the piece of the surface F the we started with (in Figure 7) and replacing it with this one.

Figures 7 through 11

The resulting surface has higher (i.e., closer to 1) Euler characteristic (Ex: how do you see that?); but, just as important, it too is a surface obtained by running Seifert's algorithm on an alternating link projection! This can be see by seeing that the half-twists that we see in every other 'end' of the new piece have the same sense as the ones in our original surface, so its boundary is still alternating. Also, this piece itself has actually become a Seifert disk for the new link - this can be seen since the orientations along its boundary are all consistent with one another. Basically, what has happened is that every other Seifert disk around the boundary has been amalgamated into one (pushing off new small disks along the other edges) - see Figure 12.

Figure 12

That completes the proof in the first case (actually, it really does! The resulting link also has non-nested Seifert circles - they're basically the same ones we started with (modulo the amalgamation).). To deal with the second case (the circles are nested), we need to introduce a new concept.

Definition: Suppose F_1, F_2 are orientable surfaces in $S^3=B^3_-\cup B^3_+$, $F_1\in B^3_+$, $F_2\in B^3_-$, with $F_1\cap F_2$ = a 2n-gon D in the equatorial S^2, as depicted in Figure 13. Then $F=F_1\cup F_2$ is a surface in S^3, called the Murasugi sum (or generalized plumbing) of F_s and F_2.

Figure 13

If we (normally) orient both F_s and F_2 to point upwards on their common disk D, then together they give a normal orientation for F. This in turn gives a sutured manifold structure to $M=S^3\setminus \text{int } N(F)$. We will prove our theorem in our second case by showing that the Seifert surface is the Murasugi sum of two surfaces, each of which is disk decomposable. Then the following result will finish our proof:

Proposition: If F_s and F_2 are both disk-decomposable, then so is F.

Proof: What we do is find a disk to decompose M along, giving a new sutured manifold which in fact is (homeomorphic to) the disjoint union of $S^3\setminus N(F_1)$ and $S^3\setminus N(F_2)$. Since by hypothesis each of these can be decomposed along disks to a union of $D^2\times$i's, and therefore M can be, too.

All we need is to find that one decomposing disk. But if we think of our picture of S^3 above as being $S^2 \times R \cup \pm\infty$, then there is an obvious disk to try - namely the complementary disk D' to the disk D in the equatorial sphere (see Figure 13). All we have to do is see that it works! But we have (by having an orientation on the two pieces F_s and F_2) already given an orientation to the boundar of (D, hence) this disk D'. Even though we have the option of choosing one of two orientations for this disk, we would be fools to simply ignore the one that the setup hands us, so we might as well see what happens if we decompose using that orientation. This is shown in Figure 14; basically, the two pieces lift off of one another. It only remains to figure out which way the two pairs of vertical sutures are joined together by the mew pieces of suture that come from D'; but it is easy to see that they basically just give us our original ∂F_s and ∂F_2 back again. Finally, by isotoping D' in each piece around (not at the same time! Pretend we have only one piece embedded in S^3 at a time) to lie parallel to D, we easily see that each piece is homeomorphic to $S^3\setminus$(the appropriate thickened surface), as desired (Figure 15).

Figures 14 Figure 16

Next time we will show how this result allows us to finish our proof for the case of alternating links. Then we will look at some more classes of knots and links for which we can fairly easily find disk decomposable spanning surfaces.

Foliations and the Topology of 3-manifolds

Outline of class 25

Today's lecture will be mostly by picture, too.

We've seen how to prove Gabai's theorem about alternating links in the case that the Seifert circles are not nested. Today we finish the proof by considering the case of nested circles.

If we (necessarily) have a collection of nested circles, then there is one circle γ which separates the rest of the circles into two (non-empty) families, those inside of γ and those outside (Figure 1). Because our link L is alternating, it is easy to see that the half twisted bands joining to γ (well, to the Seifert disk D that it bounds) from the inside all twist in the same direction (say, positive) as one another, and all the ones joining from the outside also twist in the same direction from one another, but opposite to the ones on the inside. For if two adjacent bands one inside and one outside, twisted the same way, our link would not be alternating; the same is true if two adjacent bands on the same side had opposite twisting (see Figure 2). Then by arguing inductively around γ, starting with one twisted band, it is easy to see that the bands behave as described.

Figure 1 Figure 2

Now if we imagine putting D in the plane of our projection, and lifting all of our Seifert surface F inside of γ up above the plane, while pushing all of F outside of F down below the plane, then by looking at this new picture on the side (Figure 3) it is easy to see that F is the Murasugi sum of two surfaces F_1,F_2 along (a disk very close to) D. By disk decomposing F along the complementary disk to D, using as our boundary orientation the natural one induced from ∂F, we see that the resulting two surfaces (Figure 4) are the two surfaces we would get by using the Seifert circles γ and, alternatively, the circles inside and outside of γ. In other words, they are each the result of running Seifert's algorithm on some link. But these links, it is easy to see, are each alternating (and have fewer crossings than our original link L) - this is immediate anywhere away from the loop γ (because away from γ all of the crossings are the same as for L, hence alternate)while around γ, since the half-twisted bands around the inside (resp. outside) are all of the same sign, it is easy to see that they also alternate (Figure 5). Therefore, by induction, we may assume that both of the surfaces F_1,F_2 are disk decomposable, so by the result we proved at the end of last time, F is also disk decomposable. This finishes our proof.

Figure 3 through Figure 5

This gives us a large class of knots and links for which we know how to build disk decomposable (hence norm-minimizing) surfaces. Next we will explore another class of links for which, in some cases, it is easy to identify a disk decomposable Seifert surface.

Definition: A **pretzel link** P=P(n_1,...n_k) is a link obtained by stringing k twisted bands (each given n_i half-twists) together as in Figure 6. If the integer n_i is positive, we give positive twists; if negative, we give negative twists (Figure 7). It is not hard to see that P is not a knot if two or more of the n_i are even; this is because how many components the link has is determined not by the integers n_i but by their parity (Figure 8), so we could imagine using 0 or 1-twisted bands, to count components, which makes the assertion obvious. If all n_i are odd, then P is a knot if k is odd, otherwise it has 2 components. Notice also that the ordering n_1,...,n_k is in part illusory - it should actually be thought of as a cyclic ordering, since we could instead draw our standard picture around a circle (Figure 8).

Figure 6 through Figure 8

If all of the n_i are odd, then there is an obvious spanning surface F(n_1,...,n_k) for this 1- or 2-component link, obtained by actually putting in twisted bands with an odd number of half-twists, together with the obvious disks at top and bottom (Figure 9). This surface is orientable; this is the point to having all n_i odd. By choosing a (normal) orientation on the top disk and is opposite on the bottom, the odd number of half twists allow us to pass these orientations up and down the bands in a consistent way (Figure 10). It is easy to calculate the Euler characteristic of this surface: it basically consists of two vertices (the disks at top and bottom) and k edges (the half-twisted bands), so its Euler χ is 2-k. We can use this surface to impose an orientation on P; if P has 2 components, this chooses one of the two possible orientation pairs. Then it turns out that this surface is almost always disk decomposable (hence norm-minimizing):

Figure 9 Figure 10

Theorem (Gabai): The surface $F(n_1,\ldots,n_k)$ is disk decomposable, **unless**

(1) $\{-1,1\}\subseteq\{n_1,\ldots,n_k\}$, or

(2) $(n_1,\ldots,n_k)=(n,-n)$ for some integer n.

It is easy to see that in each of these cases the resulting surface is not disk decomposable; in the first case (Figure 11), we can see that $P(n_1,\ldots,n_k) =$
$P(n_1,\ldots,n_{i-1},1,n_{i+1},\ldots,n_{j-1},-1.n_{i+1},\ldots,n_k)=P(n_1,\ldots,n_{i-1},n_{i+1},\ldots,n_{j-1},n_{i+1},\ldots,n_k)$ (Figure 12) which can be spanned by a surface with (2) higher Euler characteristic. In the second case, P is the unlink (Figure 13), and the surface F is a flat annulus, which is not disk decomposable (the resulting sutured manifold looks like Figure 14; this is an example of a norm-minimizing surface which is not disk decomposable).

Figure 11 through Figure 14

The proof, as usual goes by induction. What we do is disk decompose along the 'obvious' disks - the ones in between pairs of adjacent half-twisted bands. All that we need to do is be careful enough that in the process of the induction we never hit the two cases that we know we can't handle. Since in the first one of the bad cases the possibility of both 1 and -1 arise, our induction will be built around that.

Case 1: One of the $n_i=\pm1$.

By cyclically permuting, we may assume that $n_1=\pm1$, and WOLOG, we may assume $n_1=1$ (the proof for $n_1=-1$ is entirely similar). Then by assumption -1 does not appear in the sequence of n_i's, so in particular $n_2\neq -1$. Then, depending on whether $n_2>0$ or <0, we do one of the two disk decompositions (on the 'hole' between the first and second bands) pictured in Figure 15. This gives us, starting with the sutured manifold from $F(1,n_2,\ldots,n_k)$, a new sutured manifold, which, it is easy to see, comes from the surface $F(1,n_3,\ldots n_k)$. Since this collection of integers also satisfies our hypotheses (no -1 in $n_3,\ldots,n_k$), this surface is disk decomposable, so we are done by induction.

Figure 15

Case 2: All of the $|n_i|\geq3$, and, for some i, $n_i{\cdot}n_{i+1}>0$ (i.e., they both have the same sign).

Here again, we mean this cyclically, so the possibility $n_k{\cdot}n_1>0$ is included. We will do the case that both are positive; the case both are negative is, again, similar. Here again we do a disk decomposition along the 'obvious' disk (Figure 16); using the right orientation, doing the disk decomposition gives us a new sutured manifold, which comes from the surface $F(n_1,\ldots,n_{i-1},1,n_{i+2},\ldots,n_k)$. Then by Case 1, we are done; this new pretzel link has a 1 and no -1!

Figure 16

Case 3: No n_i is ± 1, and they all alternate sign.

Then, since they in fact alternate sign cyclically, there are an even number of twisted bands, which we can match up as positive-negative adjacent pairs (Figure 17). In this case we argue more by brute force than by induction. By doing the obvious disk decomposition between each element of a pair, depending on which orientation we use on the disk we get one of two pictures (Figure 18). Notice how the orientation of the resulting core curve differs in the two pictures. These decompositions result in sutured manifolds which do not come from surfaces, which is why we must abandon our induction. But by doing these disk decompositions for all of the pairs, using one of the orientations on one of the disks and the other on **all of the rest**, we arrive at the sutured manifold in Figure 19, which by pulling the two curves on top and bottom can be made to look like Figure 20. Now by doing disk decompositions on the remaining 'obvious' disks (from left to right, given the choices we've made thus far), we can amalgamate the single belt curves, in all of the left-hand tubes, to one (Figure 21), which swings around to give us the picture in Figure 22. Then a final disk decomposition (again, along the 'obvious' disk) gives us a 3-ball with one suture, which was our goal; we have built a disk decomposition. This completes our proof.

Figure 17 through Figure 22

Next time we will do a few more examples of these constructions, and then start our proof that disk decompositions allow us to build Reebless foliations around our Seifert surfaces.

Foliations and the Topology of 3-manifolds

Outline of class 26

We'll do one last example of a disk decomposable surface - this is the surface we get by running Seifert's algorithm on the (-2,3,7) pretzel knot (Figure 1, 2,3). The associated Seifert circles are nested, so the resulting surface is a Murasugi sum, suggesting that we first try decomposing along the disk complementary to the summing disk (Figure 4). The resulting two surfaces, it turns out, are both surfaces one would get by running Seifert's algorithm on an alternating link (the (2,2) and (2,10) torus links, respectively - see Figure 5). Since by our previous result, these are both disk decomposable, we can conclude that our Seifert surface for P(-2,3,7) is disk decomposable, hence of minimal genus. Notice that the same construction will work for any pretzel knot of the form P(2k,p,q) where p and q are both odd and have the same sign (where do things go wrong if they have opposite sign?).

Now that we know that we can find disk decomposable surfaces for lots of knots and links, it's time to convince ourselves that they really are worth having. This fact is contained in the next two theorems.

Theorem (Gabai): If (M,γ) is a disk decomposable sutured manifold, then M admits a taut (transversely-orientable) foliation $\mathcal{F}$ which is transverse to γ , and contains $R_+(\gamma)$ and $R_-(\gamma)$ as leaves. In addition, for any component N of M with $\pi_1(N)\neq 1$, $\mathcal{F}'=\mathcal{F}|_N$ has only $R_\pm \cap N$ as compact leaves, and the space of leaves of $\mathcal{F}'\setminus$(compact leaves) is S^1.

Note that M, being disk decomposable, must be a handlebody; this is because by cutting M open along a collection of disks we get a bunch of 3-balls. This means that we can obtain M by starting with a bunch of 3-balls and glue disks in their boundaries together in pairs. This amount to attaching 1-handles to the 3-balls, so M is a collection of 3-balls with 1-handles attached, i.e., a handlebody. So the proviso that $\pi_1(N)\neq 1$ is just ruling out the possibility that N is a 3-ball (and therefore, as a sutured manifold, is $(D^2\times I,\partial D^2\times I)$). Therefore no can only be foliated (transverse to the sutures) by parallel disks - this is basically Reeb stability. Note that the parenthetical comment above is actually not entirely trivial; basically (employing induction) it says that if we disk decompose a 3-ball N and get two product 3-balls $(D^2\times I,\partial D^2\times I)$, then the N itself is a product 3-ball. We leave this as an exercise (draw a picture!). This extra condition, that after removing compact leaves the leaf space of the resulting foliated (open) manifold is a bunch of S^1's, will turn out to be the only relatively hard part of the proof. We do this extra work for a reason, however, namely:

Theorem (Gabai): If (M,γ) is a sutured manifold with γ = a single annulus, and $\mathcal{F}$ is a foliation as described above, then $\mathcal{F}\cap\gamma$ = a collection of parallel simple closed curves.

Corollary: If $K\subseteq S^3$ is a knot and $F\subseteq S^3\setminus \text{int } N(K) = M$ is a disk decomposable Seifert surface for K, then there exists a taut foliation $\mathcal{F}$ of M with F as a leaf and $\mathcal{F}\cap\partial M$ is a collection of parallel longitudinal (also known as slope-0) circles.

The corollary follows immediately from the two theorems and the definition of a disk decomposable surface. Now if we were to do slope 0 Dehn surgery (also known as 0-frame surgery) on K, i.e., glue a solid torus to M so that the longitudinal circles bound meridional disks in the solid torus, then we get a new manifold M^+. By gluing meridional disks to the leaves of $\mathcal{F}$ along $\mathcal{F}\cap\partial M$, it is easy to see that we can build a foliation $\mathcal{F}^+$ of M^+, as well. But it is easy to see that this new foliation is taut, because $\mathcal{F}$ is: $\mathcal{F}^+$ has no new leaves, so the transverse loops that proved that /cafsp was taut suffice for /cafpl, as well. Notice that $F\cup D^2$ is a laef of $\mathcal{F}^+$.

But then the topological results that we proved before can come into play. In particular, Rosenberg's theorem tells us that either M^+ is irreducible, or $M^+=S^2\times S^1$, <u>and</u> /cafplsp is the standard foliation by 2-spheres. But in the second case, this would imply that $F\cup D^2$ is a 2-sphere, implying that F is a disk, and therefore that the knot K is the unknot. so we get:

Theorem (Gabai): If K$\neq$unknot has a disk decomposable spanning surface F, then the manifold obtained by doing 0-frame surgery to S^3 along K is irreducible. In addition, the surface F^+ obtained by gluing a disk in M^+ to F is (the leaf of a taut foliation, hence) is incompressible in M^+. In particular, $M^+\neq S^2\times S^1$.

A knot for which 0-frame surgery does <u>not</u> give $S^2\times S^1$ is said to satisfy Property R (from a famous list of unsolved problems in knot theory compiled by Robion Kirby). We therefore will have proved Property R for all alternating knots, for example. The conjecture is (well, <u>was</u>) that only the unknot does not satisfy Property R. This conjecture was proved by Gabai (around 1985), using exactly the techniques we have been describing. In some sense, it's what sutured manifold theory was invented <u>for</u>!

Our next (and probably last) task is to prove these two theorems. We will start with the first one. It is proved, naturally, by induction on the length of the disk decomposition. If the length is zero, then M consists of a collection of product disks, which can be given the product foliations, which clearly satisfy the theorem (the last condition is vacuous). This leaves us with the induction step - we assume we have a decomposing disk D turning (M,γ) into the sutured manifold (M_1,γ_1), and we already have a foliation $\mathcal{F}_1$ on M_1 satisfying the conditions of the theorem. What we need to do is find out how to use this foliation to find a foliation after we <u>reglue</u> the two copies of D in ∂M_1back together again (Figure 6).

We will start with a somewhat simpler case, namely the case that $D\cap s(\gamma)$ = 2 points (Figure 7), which nonetheless contains most of the idea that we will need in the more general case. This proof must necessarily break down into several cases (due in large part to the fact that the foliation on simply-connected pieces looks qualitatively different):

(1) D does <u>not</u> separate M, and $\pi_1(M_1)\neq 1$

(2) D does <u>not</u> separate M, and $\pi_1(M_1)=1$

(3) D <u>does</u> separate M, into components N_1 and N_2, and $\pi_1(N_i)\neq 1$ for each i=1,2

(4) D <u>does</u> separate M, into components N_1 and N_2, and $\pi_1(N_i)=1$ for some i (say, i=1)

We will do cases (2) and (4) first. In case (2), we have M_1 = a product 3-ball, so , since D intersects γ in 2 points, M is a product solid torus (Figure 8), i.e., an annulus$\times$I. We can therefore foliate M by parallel annuli; but this would violate our extra assumption - this would not have leaf space S^1. But it is easy to see how to alter this product foliation to make it have leaf space S^1; just cut along our disk D and reglue in a way that preserves the foliation by parallel lines, but shifts all of them 'up' (Figure 9). In case (4) we need not be so careful; D looks like a rectangle (Figure 10), so gluing N_1 to N_2 gives us something homeomorphic to N_2 (Figure 11); in fact (M,γ) is really just (N_2,γ_2), so we just use <u>its</u> foliation as the foliation of M.

In the remaining cases, we do essentially the same thing, the two copies of D are foliated (by F_1) as rectangles (Figure 11), and we just glue them together in a way that preserves the horizontal arcs of the foliation. This gives us alot of latitude in how this gluing is done; the point is that by being slightly careful about how we glue, we can insure that the space of leaves is as we describe in the theorem (provided it was before you glued).

The idea is that since the space of leaves is S^1 (when we remove the compact leaves $R_\pm$), if we look at how the vertical arc in each copy of D intersects the leaves of F_1, we see that it does so in a special way. Basically, by mapping the interior of this arc down to the space of leaves S^1, it is a locally injective map (this is just saying that because the space of leaves is S^1, upstairs a short enough arc through a leaf L intersects L is only one point. In the terminology we introduced at the beginning of the semester, all of the leaves are proper), and therefore a covering map - it does not approach a limit as we approach the ends of the arc, since if it did, then limiting point downstairs would correspond to the leaf upstairs that the arc is limiting on, i.e., $R_\pm$, a contradiction. So the transverse arc upstairs keeps hitting the same leaves in a recurring pattern as we walk form one end of the arc to the other; we can write the interior of the arc as a bi-infinite union of sub-arcs, whose endpoints are the inverse image of a basepoint downstairs, and which each map down to an arc running around S^1 exactly once (Figure 12). If we glue the two copies of D together so that they in some way respect this partitioning of the vertical arcs into subarcs, then we will be able to conclude that the resulting foliation has space of leaves S^1. To be more precise, in case (1), the two vertical arcs are being mapped into the <u>same</u> space of leaves S^1, so use the same basepoint to create the partition of the vertical arcs. In case (3), the two arcs map to different spaces of leaves S^1; choose any partition you like for each.

Then, in case (1), if we glue the two disks together so that the gluing respects the partition, in the sense that each subarc is glued to itself in what essentially the identity, then each leaf of $\mathcal{F}_1$ is glued only to <u>itself</u>, and therefore has no effect whatsoever on the space of leaves (so it remains S^1). By 'essentially the identity' we mean that each subarc is identified in a natural way to the leaf space S^1(where we imagine splitting along the basepoint), and we do the gluing so that, on the level of this natural identification, it is the identity - See Figure 13.

In case (2), we do essentially the same thing, but since the leaf spaces on each piece are <u>different</u> S^1's, we merely require that the identifications are 'essentially the same', i.e., they each correspond (instead of to the identity) to some fixed (though otherwise arbitrary) orientation-preserving (so we get a transversely-orientable foliation - the S^1's are oriented by the transverse orientation of $\mathcal{F}_1$) homeomorphism of the two S^1's, sending basepoint to basepoint. Then since one leaf in N_1 is consistently glued to the same leaf in N_2, it is easy to see that the space of leave of the resulting foliation is obtained by identifying the two S^1's together by our homeomorphism, i.e., is a circle.

Next time we will complete our proof, by showing how to carry out a similar construction in the case that D hits the sutures more than two times.

Foliations and the Topology of 3-manifolds

Outline of class 27

Today we will finish our proof that disk decomposability implies the existence of a taut foliation of the type described in the theorem. Our proof is by induction: we assume that after disk decomposition $(M,\gamma) \overset{D}{\Rightarrow} (M_1,\gamma_1)$ we have a foliation $\mathcal{F}_1$ on M_1; our task is to figure out how to build a foliation on M using this data. Last time we handled the case when $D\cap\gamma$ consists of two points; today we handle the rest of the cases.

The proof naturally breaks down into several cases (although not naturally into the order we are about to write them! - this is the order in which they are most naturally <u>proved</u>):

(1) D disconnects M and both pieces of $N_1 \cup N_2 = M_1$ have $\pi_1(N_i)=1$,

(2) D does <u>not</u> disconnect M, and $\pi_1(M_1)=1$,

(3) D does not disconnect M and $\pi_1(M_1)\neq 1$,

(4) D disconnects M, and $\pi_1(N_i)\neq 1$ for both i=1,2, and

(5) D disconnects M, and $\pi_1(\mathrm{eni})=1$ for <u>exactly</u> one i (say i=1).

Case 1: Since $\pi_1(N_i)$ are both 1 and both handlebodies, each is a disk, and so M is a disk. all that we need to verify in this case is that the set of sutures γ for M is connected, since then (M,γ) will have the structure of a product 3-ball, which we can clearly foliate as in the theorem. To prove this we just argue by induction, using as our inductive step the claim that if we decompose a 3-ball sutured manifold and get two product balls, then we started with a product ball. This is because if we have cut along a decomposing disk and end up with connected sutures, then the arcs that the decomposing disk cut γ up into before decomposing must look like Figure 1. This is because <u>after</u> decomposing the set of sutures in each piece is connected; but the arcs emanating from near the disk D must be connected to one another head to toe (Figure 2), and if anyone is <u>not</u> connected to its neighbor (Figure 3), each of the collections of arcs it separates (into two groups) have no option but to connect to one another, implying that the sutures are not connected. But then we clearly obtain a connected suture for M when we reglue.

Case 2: In this case, we have that M_1 is a product disk (by Case 1, above); our claim in this case is that the sutures of M must consist of parallel loops which are each non-trivial in $\pi_1(M)$. We prove this by (of course) induction on the number of points of $D \cap s(\gamma)$. $s(\gamma)$ comes from the suture of M_1 (Figure 4) by the usual gluing (Figure 1); if we look at the arcs of γ_1 running between the arcs parallel to the (two copies of) ∂D, each either runs from one circle to the other (non-trivial arcs), or runs from one circle to itself (trivial arcs). If it is trivial, it must go to an adjacent point (as in case 1); otherwise the argument above again implies that γ_1 is not connected. If all of the arcs are non-trivial, then it is easy to see that the resulting sutures γ of M each intersect ∂D in the same sense (i.e., always running in the same direction - see Figure 5), since otherwise it would have to turn around (creating a trivial arc). So each has non-zero intersection number with ∂D, hence are non-trivial (and hence all parallel) in $\pi_1(M)$.

If some arc is trivial, then it is easy to see (Figure 6) that the resulting sutures $s(\gamma)$ are the same as what we would obtain by using a collection of (isotopic) sutures in M_1 which give sutures in M meeting ∂D in fewer points. By induction (the base case is $\partial D \cap \gamma = 2$ points, which we handled last time), the resulting sutures (isotopic to our original ones!) are parallel and essential.

Using this, we can now build a foliation on M satisfying the theorem. We start with the standard foliation of the solid torus M by meridian disks. But we have a collection of loops (well, annuli, really (their neighborhoods)) running transverse to the meridian loops in the boundary, and cutting ∂M into two collections of annuli R_+, R_- (Figure 7). If we imagine the meridian disks as oriented upwards, then as we approach R_+, we imagine ourselves starting to pull the disks down, and as we approach R_- we pull them up (Figure 8). This gives us a (transversely-oriented, taut) foliation of M which points outward along R_+, inward along R_-, and is transverse to γ (Figure 9). Since it came from the standard foliation, it is easy to see that when we remove the compact leaves $R_\pm$, we have space of leaves S^1.

Case 3: This case is the heart of the proof; the remaining cases will amount to slight modifications in the construction we give here. We start with our disk decomposition (our pictures will illustrate the case of 4 points, for the most part), and we know that after cutting open and altering the sutures (Figure 10), we have a foliation $\mathcal{F}_1$ on M_1, with space of leaves S^1. Our task is to figure out how to use this to build our foliation $\mathcal{F}$ on M.

The basic construction is contained in Figure 11. What we do is, at the two copies of D that get glued together, we cut out what amounts to a disk$\times$I, running 'halfway' down the sutures on either side (here the disk looks like a rectangle - if we were dealing with $D \cap \gamma = 2n$ points, this would be a 2n-gon whose sides are alternately in the boundary of the annulus neighborhoods of the sutures (Figure 12)). We assume that the leaf lying at the 'bottom' of the cut is the same leaf L on both sides (by cutting down that far - the leaves on the left are the <u>same</u> as the leaves on the right (since after throwing away the compact leaves we have space of leaves S^1, so the picture along transverse arcs are the same on both sides, cycling through the same set of leaves infinitely often). Did that last sentence make any sense? What we will do is glue these two disks at the bottom of the cuts together. Because the normal vector to the foliation $\mathcal{F}_1$ is pointing outward at the bottom of one cut and inward on the other, we will get a transversely-oriented foliation after gluing. Now in the sutures, beyond the point where we cut down to, there is another rectangle in the suture, with one end on $\mathrm{R}_\pm$ (basically, the 'other half' of the suture). We will <u>also</u> glue the walls of the cuts we have made to the corresponding rectangles given in these other halves; see Figure 13. Provided we glue these pairs of rectangles together by homeomorphisms that preserve their foliations (from $\mathcal{F}_1$) by horizontal arcs, it is fairly easy to see that we will succeed in building a foliation of M which has $\mathrm{R}_\pm(\gamma)$ as leaves and be transverse to the sutures (Figure 13). We have a very wide latitude in how these gluing take place; as we did last time, we will see that by choosing these gluing carefully, we can <u>also</u> insure that the space of leaves <u>remains</u> S^1.

If we take the arc transverse to each of the rectangles that we are gluing in pairs, it easy to see that each starts at L and then runs transverse to $\mathcal{F}_1$ tending to $\mathrm{R}_\pm(\gamma_1)$. This path runs either with or against the normal orientation of $\mathcal{F}_1$, depending upon which side of the picture the rectangle originates (i.e., lies in the cut) from. Because the space of leaves (when we throw away the $\mathrm{R}_\pm$) is S^1, if we use (the point corresponding to) L as basepoint, we can partition the intervals, as before, into an infinite number of subintervals all wrapping around S^1 once as we march out to $\mathrm{R}_\pm$ (Figure 14). If we do our gluings, as before, to respect this partitioning (i.e., after the identification with S^1, are gluings over each subarc is actually the identity), then we will always be gluing a leaf of $\mathcal{F}_1$ to <u>itself</u> (again, as in the corresponding case from last time). Therefore, $\mathcal{F}$ will have the <u>same</u> space of leaves as $\mathcal{F}_1$ (essentially, it has the same <u>leaves</u>, after throwing away $\mathrm{R}_\pm$ (there could by components being glued to one another in $\mathrm{R}_\pm$)). This proves the inductive step in this case.

In **Case 4**, things are essentially the same as before; the leaves on one side are all different those on the other, so we just choose any two leaves L, L$'$ to cut down to (Figure 15). Then the same structure occurs on the transverse arcs on either side (but measured against <u>different</u> spaces of leaves S^1), and by choosing a single basepoint-preserving, orientation-preserving homeomorphism of the spaces of leaves, and making our gluings to respect those identifications (as we did last time), we again can achieve space of leaves S^1.

Finally, in **Case 5**, we can argue one of two ways. The foliation on the left-hand side (say) is a product foliation by disks; and we are gluing this to one with the usual space of leaves foliation. The transverse arcs on the left hand side intersect the same leaves (on the inside of the cut; on the outside, they all intersect the same (complementary) set of leaves). If we make one each of the two types of gluings we need, and use it to create an (unnatural) partitioning of these disks into subarcs, and then make all of our gluings respect that identification, we will again achieve space of leaves S^1. Alternatively, one can just notice that in fact (M,γ) and (N_2,$N_2 \cap \gamma_1$) are homeomorphic as pairs, so we could just use the old foliation; we leave this as an exercise.

This completes our construction of taut foliations from disk decompositions. We leave it as an exercise to show that these foliations are (inductively) taut (imagine starting with a solid torus). This leaves us with one last task - to show that if we build such a foliation around a Seifert surface for a knot, the foliation intersects ∂N(K) in parallel circles. This will be our task for next time; in the proof we will need to know the following fact (which we could have prove long ago, but who knew we'd need it?!).

Proposition: If M is a compact manifold with a (transversely-oriented) codimension-1 foliation $\mathcal{F}$, then the set C of all compact leaves of $\mathcal{F}$ is a closed subset of M.

This is actually true in all dimensions - we will give the general proof. The hypothesis of codimension-1 is essential, as is the compactness of M (try, for example, to build a counterexample using a foliation of $R^2 \setminus \{pt.\}$by circles and lines). To prove it we will first need the following fact:

Lemma: If M, $\mathcal{F}$ are as above, and $\mathcal{U}_1, \ldots \mathcal{U}_n$ are distinguished coordinate charts for $\mathcal{F}$, then there is a number N such that for any compact leaf L of $\mathcal{F}$, L meets each $\mathcal{U}_i$ in at most N plaques.

Proof: Suppose not; suppose that for any n there is a compact leaf L_n of $\mathcal{F}$ and an i such that L_n meets $\mathcal{U}_i$ in more than n plaques. Using the principle that if you throw an infinite number of things into a finite number of holes, you must have thrown an infinite number of things into one of the holes, we can assume that i=1 for all n. Now take a transverse arc α to the distinguished chart $\mathcal{U}_1$ (see Figure 16), and look at the $L_n \cap \alpha$. This consists, for each n, of more than n points. However, since each L_n is compact, this also consists of a finite number of points for each n - otherwise, these points of intersection would give an infinite discrete set in a compact leaf, which is impossible. But now we will inductively choose a sequence of compact leaves (which we will still call L_n) with the following property:

For each L_n, there is a loop γ_n transverse to $\mathcal{F}$ such that $L_n \cap \gamma_n$= 1 point, and $L_k \cap \gamma_n = \emptyset$ for all k<n.

If we have found $L_1, \ldots, L_{n-1}$, then together these leaves intersect α in some finite number K of points. By our hypothesis, we can find a leaf L_n of $\mathcal{F}$ which meets α in more than K+2 points; by the pigeonhole principle (the K points of the previous leaves cut α into K+1 subarcs) there are two of these points x,y which have none of the points of intersection of $L_1, \ldots, L_{n-1}$ between them along α. But then by using one of our standard tricks from long ago, we take an arc in L_n between these two points, and by pushing it off of itself using the transverse orientation to give us the loop γ_n that we desire; it only hits the leaves that pass between x and y, none of which are our previous compact leaves.

This construction allows us to find a contradiction to our hypothesis. Since $L_n \cap \gamma_n = 1$ point, on the level of homology, their intersection number $I([L_n],[\gamma_n]) = \pm 1$, and, in particular, $[L_m] \neq 0$ in $H_{m-1}(M)$ (where m=the dimension of M)for every n. Also, the facts that $I([L_k],[\gamma_n]) = 0$ for every $k<n$ and intersection number is linear implies that the homology classes $[L_1],\ldots,[L_n]$ are linearly independent in $H_{m-1}(M)$ for every n (this is some fairly easy linear algebra). This in turn implies that the dimension of $H_{m-1}(M)$ is at least n, for every n! so $H_{m-1}(M)$ is infinite-dimensional, contradicting the fact that M is compact.

So we have obtained a contradiction, proving the existence of our universal bound N. Next time we will show how to use this to prove the closed-ness of the set of compact leaves, and then prove our last remaining theorem about foliations of knot complements.

Foliations and the Topology of 3-manifolds

Outline of class 28

Last time we showed that for any compact M and (transversely-oriented) foliation $\mathcal{F}$ of M, and for any finite collection of distinguished coordinate charts for M, there is a number N such that any compact leaf of $\mathcal{F}$ intersects each of our charts in at most N plaques. Today we will use this to prove:

Proposition: for M and $\mathcal{F}$ as above, the set C of compact leaves of $\mathcal{F}$ form a closed subset of M.

and then we will use this fact to prove our remaining theorem related to disk decomposable surfaces.

Proof: Suppose C is <u>not</u> closed; then there exists a sequence of points x_n in C (i.e., $x_n \in L_n$ for some com,pact leaf L_n of $\mathcal{F}$) limiting on a point $x \notin C$ (i.e., $x \in L$, where L is a non-compact leaf of $\mathcal{F}$). But since M is compact, we can cover it by finitely-many distinguished coordinate charts for $\mathcal{F}$. Since L is non-compact, for one of these charts $\mathcal{U}$, $L \cap \mathcal{U}$ consists of an infinite number of plaques (otherwise L can be written as a union of a finite number of plaques; by assuming we chose our charts so that their closures are 'closed' charts inside some larger chart, we can conclude that L is the union of finitely-many compact disks, hence is compact).

The point x is in some chart; call it $\mathcal{U}'$. If we take transverse arcs α, α'through the middle of $\mathcal{U}$ and through x, then $L \cap \alpha$consists of an infinite number of points. Also, since the x_n limit on x and are not contained in L (since L is not compact), there must be plaques in $\mathcal{U}'$ of the leaves L_n limiting transversely down on the plaque containing x (Figure 1). But we can quickly use these two facts against one another; we have a bound N on the number of times the leaves L_n hit $\mathcal{U}$. But if we pick N+1 distinct points in $L\cap\alpha$and join them by arcs in L to the point x (Figure 2), then by looking in the normal fence over each arc (i.e., imagine lifting each arc in the normal direction) we get arcs in the L_n joining the x_n to points lying over our N+1 points, for large enough n. But this implies (since our N+1 points are separated from one another by some definite distance, and we wait until our points are less than half that distance from their corresponding points of L) that there are at least N+1 distinct plaques of L_n in $\mathcal{U}$, for large enough n, a contradiction.

We can actually get more from this same line of argument, namely:

Proposition: The set of compact leaves of $\mathcal{F}$ fall into finitely-many homeomorphism classes, and if a sequence of compact leaves L_n limit on the (compact) leaf L, then for sufficiently large n $L\cong L_n$, and bound a product between them.

Proof: The second part implies the first: if $\mathcal{F}$ has compact leaves L_n which are pairwise non-homeomorphic, then by choosing a point x_n of each, we get a sequence in the compact M, so has a convergent subsequence (call it x_n) converging to some x (which by the above is contained in a compact leaf L). But the second part says that then for large n, $L\cong L_n$, so in particular are homeomorphic to one another, a contradiction.

For the second part, we simply use the fact that any compact surface, such as our leaf L, can be thought of as a 2-disk with arcs in its boundary identified (a similar argument works when the dimension of M is larger than 3 (provided you know that a compact n-manifold can be written as an n-ball with faces in its boundary identified)). This 2-disk can be lifted along the normal fence to 2-disks in all sufficiently nearby leaves; in particular, to any other compact leaf sufficiently close. But now if one of our disks lies in the compact leaf L_n, then the claim is that $L_n$is obtained by exactly the same gluings on this disk as give L (hence they are homeomorphic). This basically amounts to saying that if we draw an arc α in the disk in L joining to faces that are identified (Figure 3), then the lift of this arc to L_n joins two faces of the disk which are identified in L_n. But this follows because L_n is compact (and $\mathcal{F}$ is transversely-oriented): If we take the normal fence over α with its endpoints identified, we get an annulus, and the lift if the arc into L_n is, we claim, a loop (giving what we want). But this is easy to see, because if it weren't, then (in one direction) traversing the arc returns us to a point in the fence closer to the bottom from where we started (Figure 4). But then continuing around again brings us still closer, and continuing, will find an infinite number of points in L_n lying in the transverse arc over a pont of L. But this gives an infinite discrete set in L_n, a contradiction, since L_n is compact.

Since the gluings correspond, it is easy to see that the product structure between the two disks (in L and L_n) glues up to give an I-bundle between L an L_k, which, since $\mathcal{F}$ is transversely oriented, must be a product.

With this in hand, we can now give a proof of our last theorem. We will actually prove something slightly stronger:

Theorem: If K is a knot in S^3, $M=S^3\setminus\text{int } N(K)$, and $\mathcal{F}$ is a taut, transversely-oriented foliation of M such that:

(1) some leaf F of M is a Seifert surface for K,

(2) for every non-compact leaf L of $\mathcal{F}$, $\overline{L}$=L$\cup$compact leaves of $\mathcal{F}$, and

(3) $\mathcal{F}\cap\partial M$ is a foliation of ∂M with no Reeb annuli (Figure 5),

then $\mathcal{F}\cap\partial M$ is a foliation of ∂M by parallel (longitudinal) circles.

Proof: Suppose not: then since $\mathcal{F}\cap\partial M$ contains a compact loop (the boundary of the Seifert surface), and since we know that the set of loops is closed, we must be able to find a situation like Figure 5. We are assuming that the set of compact loops of $\mathcal{F}\cap\partial M$ is not the entire boundary, so if we start walking transverse to the foliation (starting at a compact loop) we will find a <u>last</u> closed loop; beyond that we must have a non-compact leaf λ spiralling down on a compact loop γ. Call the leaf of $\mathcal{F}$ containing λ L, and the leaf of $\mathcal{F}$ containing γ S. Notice that we immediately have $\gamma\subseteq\overline{\lambda}\subseteq\overline{\ell}$, so $S\cap\overline{L}\neq\emptyset$, and consequently $S\subseteq\overline{\ell}$ (Exercise: just take any point of S and join it to a point of γ by an arc in S). Consequently, by (2), S is compact (or S=L; we will show why this is impossible after the rest of the proof).

Now we show that $\partial S=\gamma$; this follows from our work on Thurston norms! This is because the homology class $[\partial S]$ goes to 0 under the homomorphism $H_1(\partial M)\rightarrow H_1(M)$ (since it is clearly a boundary in M!). It therefore is represented by $n\cdot\ell$, where ℓ is the longitude of our knot (this is basically the definition of longitude). But by Alexander duality, S then represents n in $H_2(M)$; but since S is connected, and we showed that if a surface represents $\alpha+\beta$, it is the union of two disjoint surfaces each representing one of the classes, it follows that $n=\pm1$. But since our hypothesis of no Reeb annuli in ∂M implies that all of the compact leaves in ∂M are parallel and coherently oriented (Figure 6), it then follows that ∂S is connected.

If we take a parallel copy γ' of γ lying on the λ-side of γ, lying close to γ, we can make $\gamma'\cap\lambda=$ a single point (Figure 7). But since clearly no compact loop of $\mathcal{F}\cap\partial M$ lies on the λ-side of γ, no compact <u>leaf</u> of $\mathcal{F}$ lies near S on the λ-side of S (otherwise, we could, by taking an arc in S from a point of S limited on by compact leaves on the λ-side to γ, we could conclude that these compact leaves intersect ∂M (in loops!) on the λ-side of γ). This means we cam imagine (by pushing S off of itself in the λ-direction) that γ is bounded by a surface S′ which is parallel to S and does not intersect any compact leaves of $\mathcal{F}$.

But since the set C of compact leaves of $\mathcal{F}$ is closed, S$'$ and C are disjoint closed subsets of M. By normality, there is then an open sets $\mathcal{U}$, $\mathcal{V}$ of M such that C$\subseteq\mathcal{U}$, S$'\subseteq\mathcal{V}$, and $\mathcal{U}\cap\mathcal{V}=\emptyset$. But because $\overline{\mathrm{L}}$=L$\cup$compact leaves, it follows that L$\setminus\mathcal{U}$ is compact <u>in</u> <u>L</u>. Otherwise, there is a sequence x_n in L$\setminus\mathcal{U}$ with no limit in L; but then since M$\setminus\mathcal{U}$ is (closed in M, hence) compact, some subsequence has a limit (in the transverse direction) in M$\setminus\mathcal{U}$). This point is in $\overline{\mathrm{L}}$, but not in C, so would have to be in L, again contradicting the statement we haven't proved yet.

Therefore, since S$'\cap\mathcal{U}=\emptyset$, we have S$'\cap$L=S$'\cap$(L$\setminus\mathcal{U}$) is the intersection of two compact surfaces. By wiggling S$'$ slightly (keeping it in $\mathcal{V}$) we can make this intersection transverse, so S$'\cap$(L$\setminus\mathcal{U}$) consists of circles and arcs properly embedded in each (since the boundary of one intersects the other only in their common boundary - this is clear for the boundary of L$\setminus\mathcal{U}$ coming from ∂M, and the <u>other</u> ∂-components of L$\setminus\mathcal{U}$ come from $\partial\mathcal{U}$, which S$'$ avoids). But this quickly causes trouble; arcs have two endpoints, but where is the other endpoint of the arc containing $\gamma'\cap\lambda$? This is the only point of ∂S$'\cap$(L$\setminus\mathcal{U}$)! This contradiction proves the theorem, once we establish:

If $\overline{\mathrm{L}}$=L$\cup$compact leaves, then L cannot limit transversely on itself. for if it did then if we look at a transverse arc α through such a limit point x, every point y of B=$\overline{\mathrm{L}}\cap\alpha$ is a limit point of B. If y$\notin$L, this is clear (y <u>is</u> a limit point, and this limiting must be from the transverse direction rather than from in the leaf containing it) while if y$\in$L, by joining it to x by an arc in L and looking at the normal fence over this arc, by the usual argument, we find arcs in L limiting down on this arc, giving, at the end y, a sequence of points in L limiting down on y. $\overline{\mathrm{L}}\cap\alpha$ is therefore what is known as prefect. Of course, $\overline{\mathrm{L}}\cap\alpha$ is also closed (in the interval α).

But a closed, perfect subset of an interval is <u>uncountable</u> (this is one of my favorite facts that I have no idea how to prove!). so if L limits on itself, then $\overline{\mathrm{L}}\cap\alpha$ is uncountable.

But this is impossible, if $\overline{\mathrm{L}}$=L$\cup$compact leaves. For L$\cup\alpha$ is countable (if it were uncountable, the surface L would contain an uncountable discrete set, which is impossible). On the other hand, ($\overline{\mathrm{L}}\setminus$L)$\cap\alpha$ must be finite; if it were infinite, since any compact leaf intersects α in finitely-many points, there must be infinitely-many such compact leaves in $\overline{\mathrm{L}}$; but then by choosing a point in each leaf (in α), we can find a sequence z_n each in a distinct compact leaf L_n, converging to some point z (which by our first result today is in a compact leaf L$'$). But then we also know that for large enough n, the L_n and L are all parallel to one another; but by choosing three such leaves (Figure 8), with products in between them, it is impossible for a single leaf L to limit on <u>all</u> of them - to limit on the middle one, it must be in one of the two products, but whichever one it's in, the third leaf cannot be limited on by L.

This finishes our proof of the theorem, which finishes our study of foliations arising from disk decompositions, which more or less finishes what I wanted to present to you in this class. I'll lecture next time on one of my favorite bizarre facts from foliation theory, that $\mathbf{R^3}$ can be foliated by circles, mostly so I can get this last set of notes into your hands. I don't plan to write that up, so this will be our last installment. I suppose these notes could use a bibliography or something, but I don't know when I will get around to it. Technically, I also owe you a proof of that thing I skipped long ago, that an immersed, null-homotopic loop in a surface can be turned into an embedded one by a sequence of Reidemeister-like moves; I'll get to it one of these days!