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 first time we encounter something other 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 not 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.

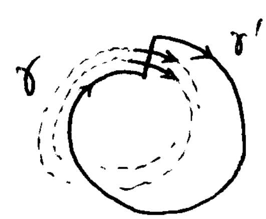

(3): We run into another center singularity. We'll deal with this case after the next two; in this case we shall see we have a 2-sphere, and have found our one exception to the rule of always finding vanishing cycles.

(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 <u>hit</u> 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:

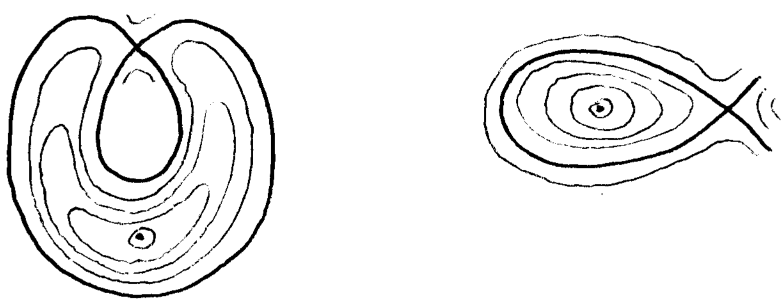

We can't have our center circles running into 3 or more of the corners of the singular leaf (because an arc joining two of the arcs emanating from the saddle forms a (separating) loop, which must therefore be joining an adjacent pair of such arcs), so there are only these two pictures (we can't run into 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 <u>not</u> null-homotopic in its leaf, we are done: it is a vanishing cycle by the argument in case (1) above. If it <u>is</u> null homotopic in its leaf, we then cancel the center and saddle singularity as in the sequence of pictures below:

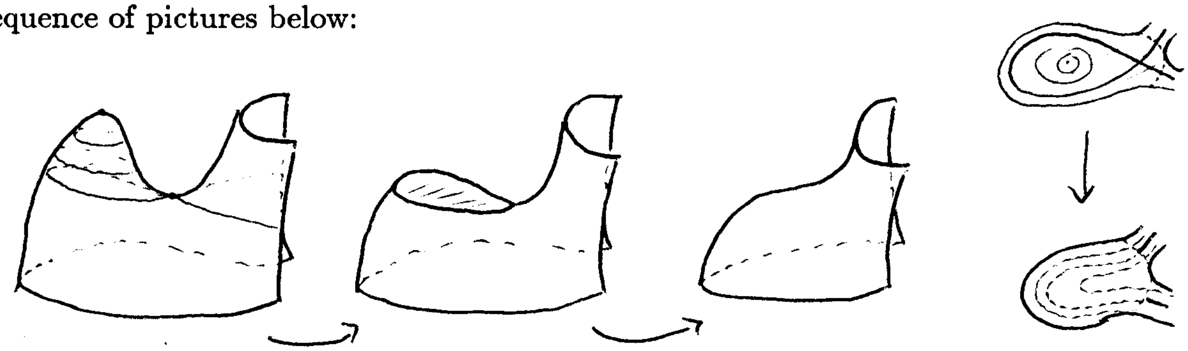

We first redefine f on the disk (containing the center) cut off by γ, so that it instead maps into the leaf containing γ (using the fact that it is null-homotopic). This map is no longer Morse; all of the points in the disk are degenerate critical points. But if we 'smooth' off the map in a neighborhood 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 $f:S^2\to M$ 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 $r_k<t_k<s_{k+1}$ and $r_k<s_k$ 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 $r_k<x_k<s_k$. Then the disk at level t_k lifts down and up to x_{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!).

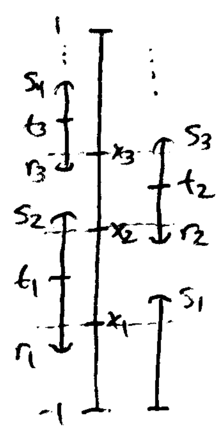

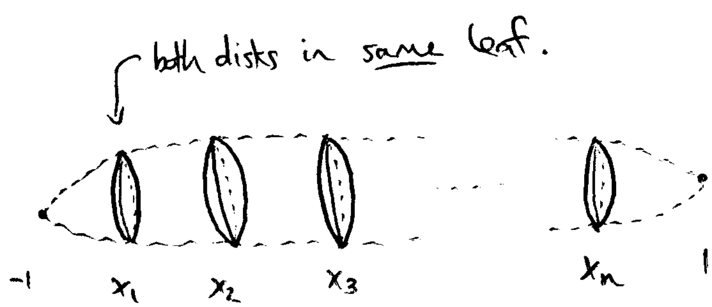

So we have now finished the first part of our outline of the proof of Novikov's theorem. We have found (except in the exceptional case $M=S^2\times S^1$) a (singular) vanishing cycle γ under each of the three hypotheses of Novikov's theorem. By general nonsense (transversality theory), since we are dealing 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 embedded; 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:

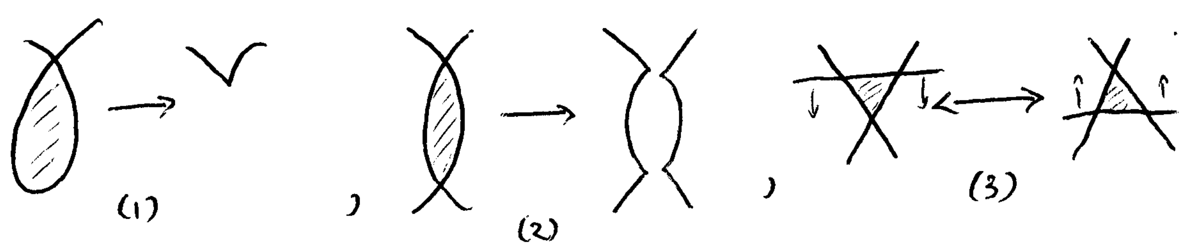

The reader familiar with knot theory will recognize that these moves are the 'immersed' versions of Reidemeister's moves, by which a knot can be transformed to any other knot isotopic to it.

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 $\gamma=\gamma_0$ is our immersed vanishing cycle, with n self-intersections, and consider the family of lifts γ_2, $0\leq t\leq\epsilon$, 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).

t's s.t. the push is null-homotopic

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.

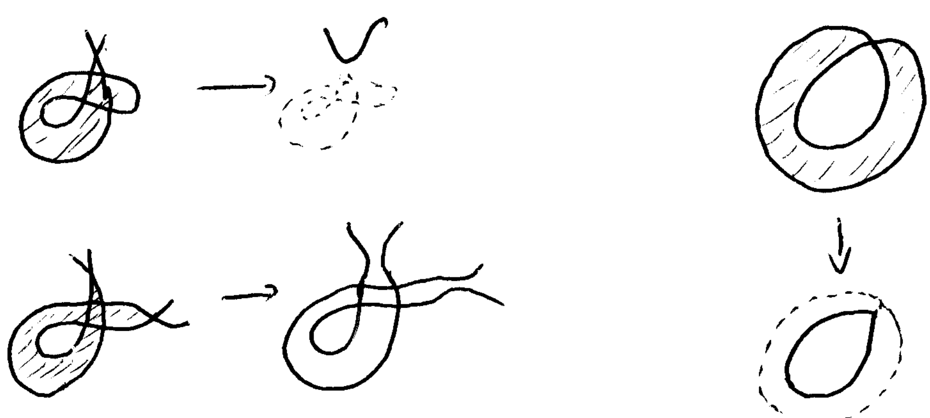

It is worth noting that we can in fact have a non-embedded vanishing cycle whose lifts are all embedded; we can build one by taking the meridian loop in the boundary of a Reeb component and drag a piece of it around the torus in the longitudinal direction until it intersects itself; see the figure.

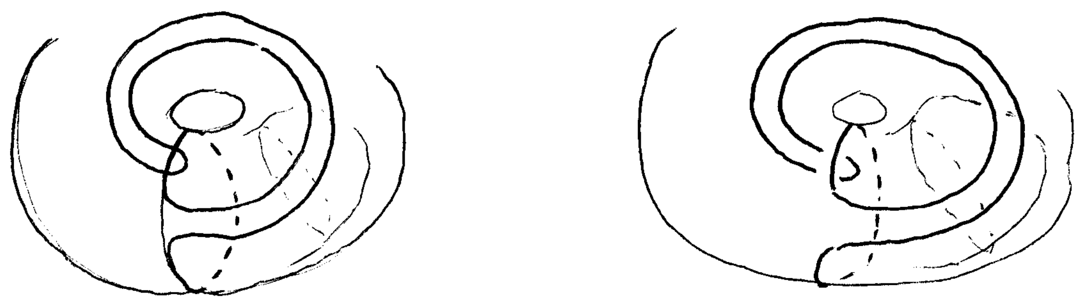

We will now use this new 'improved' vanishing cycle to find an embedded vanishing cycle. The idea is that since we now have embedded, null-homotopic, lifts, they bound embedded disks D_t in their leaves. In particular this is true of γ_ϵ; it bounds D_ϵ. What we will do is look at the intersection of 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 <u>it</u>) bounded by γ_{t_2}, $t_2<t_1$; continuing inductively, we could then in fact conclude the D_ϵ hits (in fact <u>contains</u>!) an infinite number of the γ_t's, a contradiction. We therefore have a picture like in the figure below.

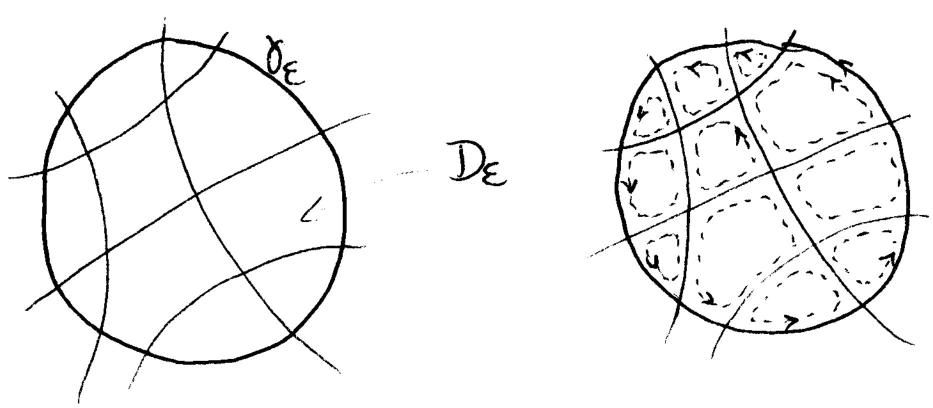

These arcs cut the disk up into pieces (subdisks, since no piece can have disconnected boundary (because there are no loops)). Each piece therefore gives us a loop in the disk. There are now two things to verify, to finish our proof:

(1): <u>If</u> all of these loops push down all the way to γ_0 as closed loops, then the image of one of them is <u>not</u> 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).

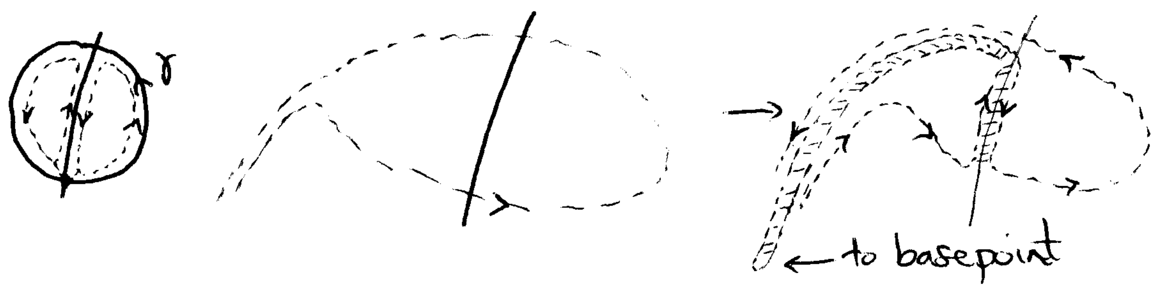

The second one is a little tricky. As a matter of fact, it is SO tricky that I don't know how to prove it. I still think it's true, but the proof I had in mind goes sour at one point (I don't know that what ought to happen does). BUT as it happens, I know of a different way to show how to pass from a 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.

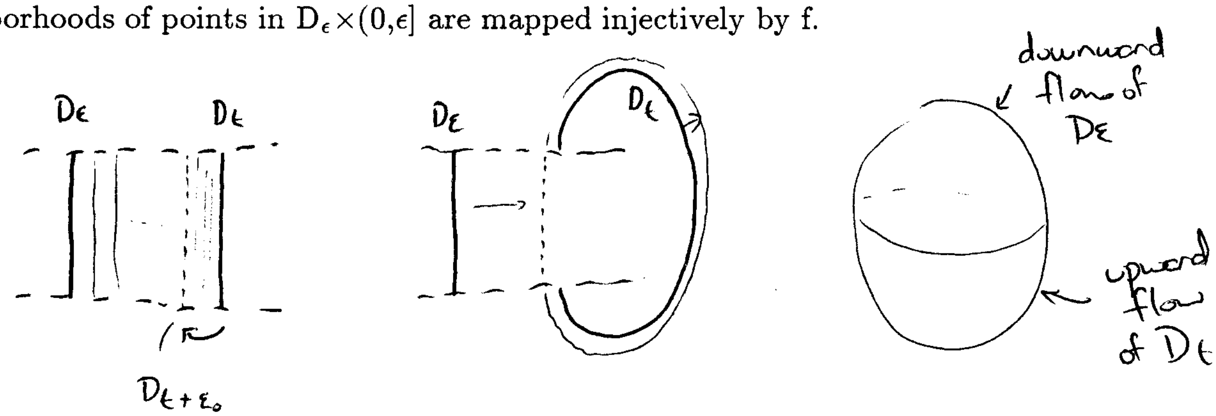

But again, this map cannot be extended to t=0; if it could, it would give a map of a disk into the leaf L, restricting to the loop γ on the boundary, implying it is null-homotopic, a contradiction. Then, again as in the embedded case, some trajectory out of a point $x \in D_\epsilon$ becomes infinitely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

Now we rejoin our embedded proof. Look at our function f restricted between these two disks; $f: D_\epsilon \times [t_0,\epsilon] \to M$. 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 $D_{t0} \setminus D_\epsilon$. $D_{t_{n-1}}$

For large enough n, the points of $D_{tn} \setminus D_\epsilon = f(A_1)$, 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 doen'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

* because the orientation of the annulus (coming from the leaf) together with the orientation of the trajectory <u>always</u> gives the orientation of the manifold M.

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 $H_2(M)$.

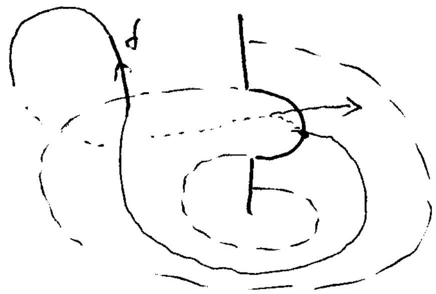

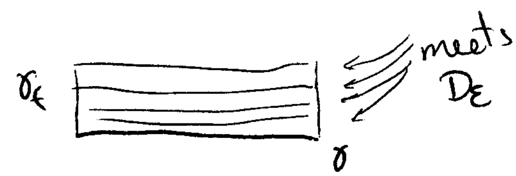

But now the annulus A_1 can be mapped, by the limiting values of the trajectories, into the leaf L containing γ; furthermore, since $f(A_1)$ is in a leaf of $\mathcal{F}$ and so transverse to the trajectories, the same proof as outlined above show that this limiting map is a local homeomorphism. Further, 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 $\mathbf{R}^2$ 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 $\mathbf{R}^2$ 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, $\gamma = \tau^{\mathrm{d}}$ in $\pi_1(\mathrm{T})$, since τ^{d} lifts in $\mathbf{R}^2$ 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 $\mathrm{f(a)<a}$, 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.

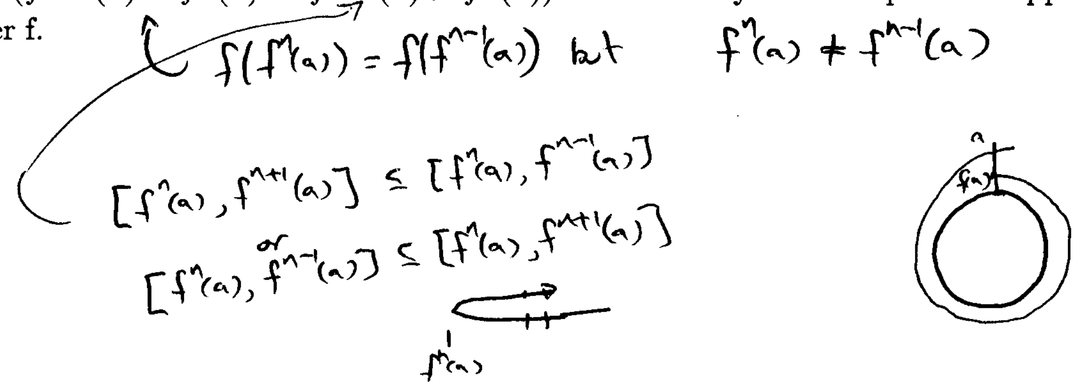

But since τ^{d} is not null-homotopic, τ is, and since the lifts of τ^{d} are loops null-homotopic in their leaves (since they are homotopic (in the leaves)to the lifts of γ - just lift the homotopy downstairs), the lifts of τ are null-homotopic. This is because the lifts of the d^{th} iterate of τ are 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 $\mathrm{M} \neq \mathrm{S}^2 \times \mathrm{S}^1$, 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 $S^2 \times S^1$, 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 $\pi_2(M) \neq 0$ 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 $\mathbf{R}^2$. Therefore if we lift $\mathcal{F}$ to a foliation $\tilde{\mathcal{F}}$ of the universal cover $\tilde{M}$ of M, then for every leaf $\tilde{L}$ of $\tilde{\mathcal{F}}$ (which projects as a covering map to a leaf L of $\mathcal{F}$), the composition (using the obvious maps) $\pi_1(\tilde{L}) \rightarrow \pi_1(\tilde{M}) \rightarrow \pi_1(M)$ is (equal to $\pi_1(\tilde{L}) \rightarrow \pi_1(L) \rightarrow \pi_1(M)$, which is) injective, so $\pi_1(\tilde{L}) \rightarrow \pi_1(\tilde{M}) = 0$ 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 $\tilde{M}$ 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 $\tilde{M}$; 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 $i:S \rightarrow M$ 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 $i^*\mathcal{F}$ 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 $i^*\mathcal{F}$ 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!

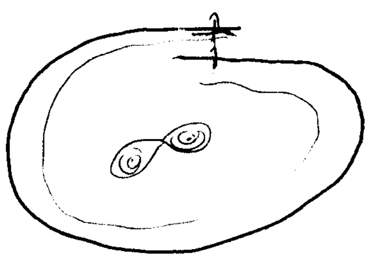

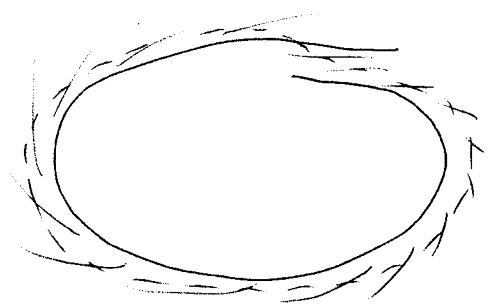

Now if we pick a center singularity of $i^*(\mathcal{F})$ and work our way out (as in the previous argument), only two things can happen; either we bump into another center singularity (so $i^*(\mathcal{F})$ has no saddles, which is what we want), or we bump into a saddle singularity (once or twice) - 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), assume 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∩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 surger 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 collection 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 cancel a center (one of the ones 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:

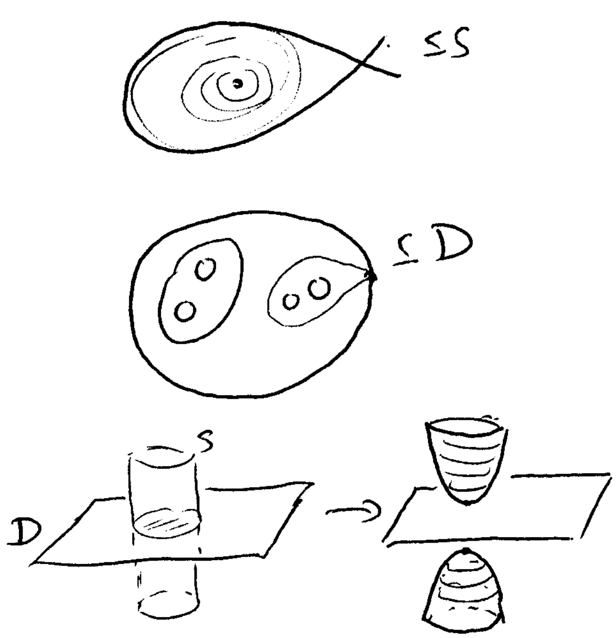

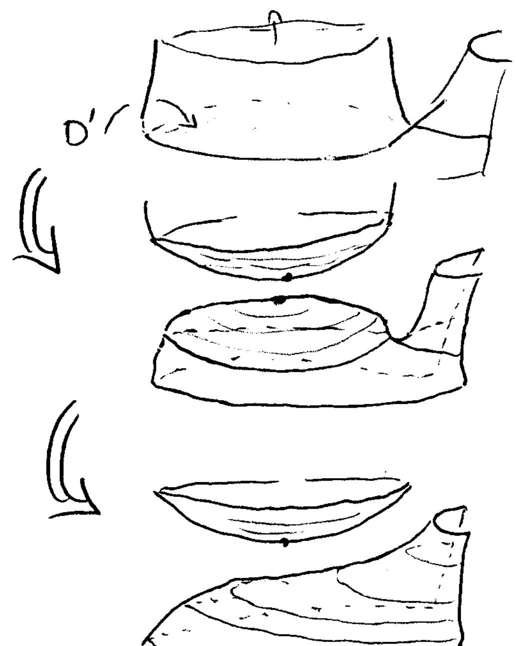

Proposition: If S⊆M is a 2-sphere and D is a disk in D with $S \cap D = \partial D = \gamma$, splitting S into disks D^+, D^-, and both of the 2-spheres $S^+ = D^+ \cup D, S^- = D^- \cup D$ bounds a 3-ball, then S bounds a 3-ball.

Proof: The proof is by picture! If both bound 3-balls (B^+, B^- respectively), then $\partial B^+ \cap \partial B^- = D$; if $\text{int}(B^+) \cap \text{int}(B^-) = \emptyset$, then it is easy to see (see figure) that $B^+ \cup B^-$ is a 3-ball, with boundary S; if $\text{int}(B^+) \cap \text{int}(B^-) \neq \emptyset$, 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 total 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 $t_1 \neq t_2$, $D_{t_1} \cap D_{t_2} = \emptyset$.

Proof: If $D_{t_1} \cap D_{t_2} \neq \emptyset$, then since $\partial D_{t_1} \cap \partial D_{t_2} = \gamma_{t_1} \cap \gamma_{t_2} = \emptyset$, 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 $D_{t_1} \subseteq D_{t_2}$. But then the arc $[t_1,t_2]$ in S together with an arc in D_{t_2} 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 $D_{t_i} \times [t_i-\epsilon, t_i+\epsilon]$. 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 $\mathbf{R}^3$ is irreducible (this is known as Alexander's Theorem), since $\mathbf{R}^3$ 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 canelled 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 1~~4~~5

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) incompressible 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 not 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 disk in F), contradicting Novikov's theorem. So we therefore run into a saddle; as before, this can happen in one of two ways (see figure). Unlike the cases we have dealt with before, we cannot avoid the second case - it is why we must allow for circle tangencies.

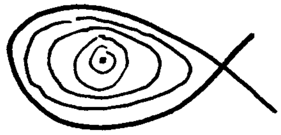

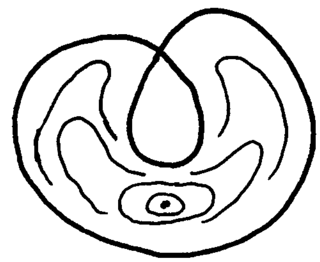

Let us deal first with the case we already understand - when we run into a saddle once when moving out from a center. This gives us an embedded loop in a leaf L of $\mathcal{F}$; since it is null-homotopic in F (hence M - it bounds a disk in F), it is therefore null-homotopic in L, hence bounds a disk 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.

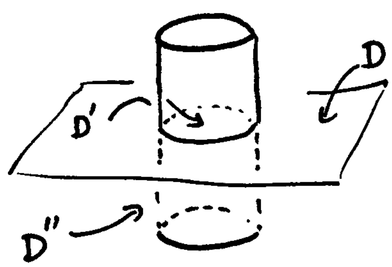

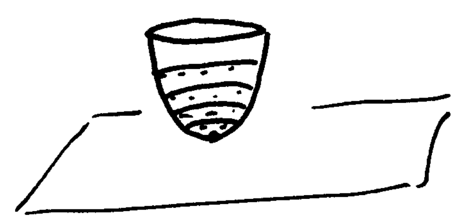

Then continuing this process we can make D$\cap$F=∂D; then as before we can cancel the center and saddle with one another, reducing the number of centers. After doing this a finite number of times, we can assume that all of the centers that are left (if any) meet their corresponding saddles twice. Notice 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~~ δ, 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.

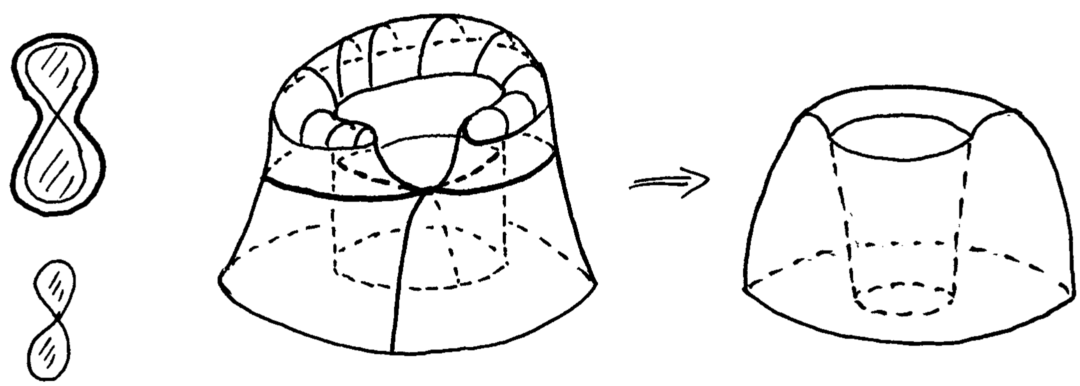

We should note that in many circumstances we can remove circle tangencies in F; if the loop representing the tangency has non-trivial holonomy around it in the singular foliation of F (see figure), we can, by pushing the tangency down a little to make the loop <u>transverse</u> to the singular foliation,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

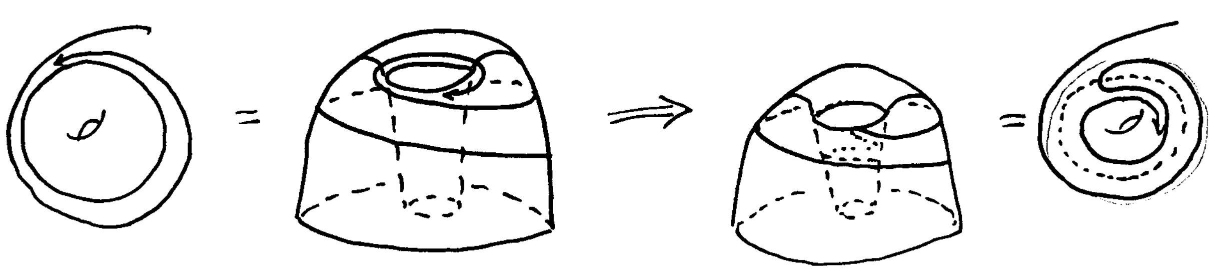

We can also improve on this result, if we assume the F is a torus and $\mathcal{F}$ is taut:

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 $\chi(T)=0$, 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.

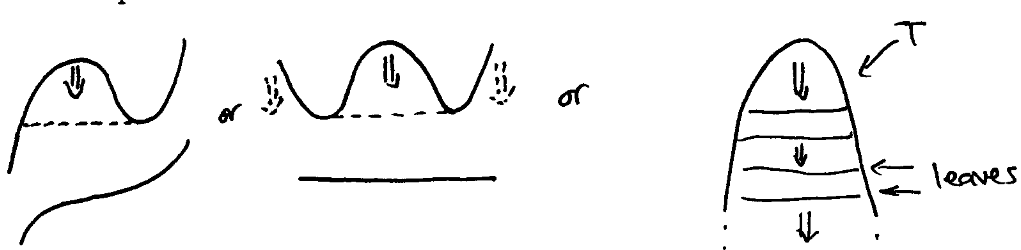

As we push the annulus A down into the leaves of $\mathcal{F}$, we build a function from $A\times[0,\infty)$ to M, and a point must trace out a trajectory that is infinitely long (otherwise, there is a limiting map of an annulus into a leaf, which we could push further). This trajectory must pile up 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 $A\times[0,t]/\sim$, 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 can 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.

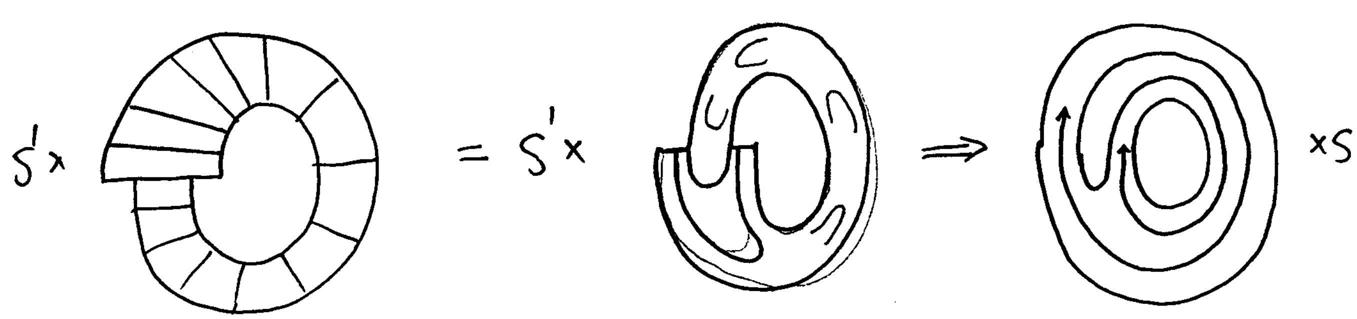

Foliations and the Topology of 3-manifolds

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 ξ 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 $e(\xi)$ of ξ 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 ξ (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 ξ, 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 $\xi|_{\sigma^2}$ is 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 $w(\sigma^2)$. Then we define $e(\xi)$ of a 2-chain $\Sigma\sigma_i$ to be $e(\xi(\Sigma\sigma_i) = \Sigma w(\sigma_i)$.

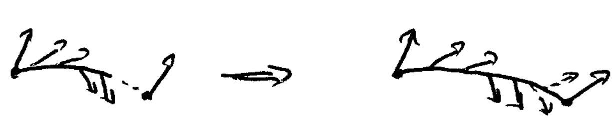

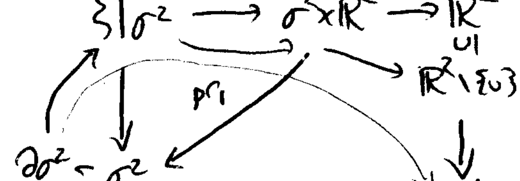

changing this section (by, for example, adding several windings along one of the 1-simplices - see figure) will certainly change the winding numbers, so will change the value of the Euler class. The point, however, is that for 2-cycles, the value of the sum is independent of these choices. This is because 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)

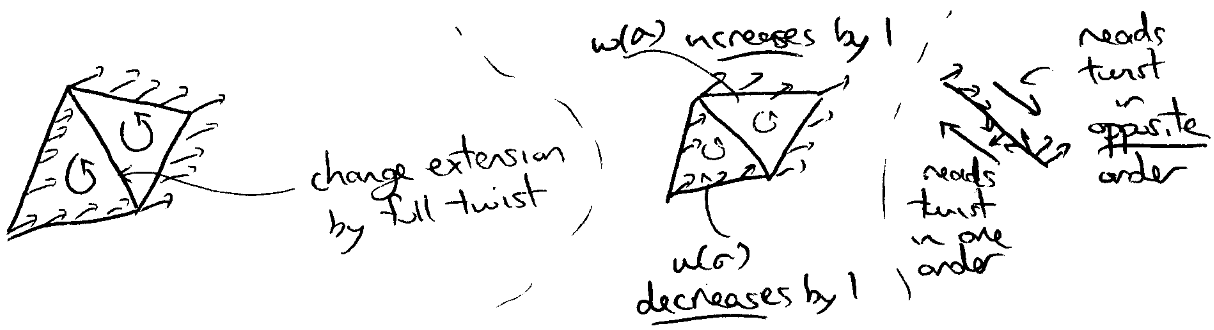

This at least gives us a function from 2-cycles to **Z**. With a little bit more work, you can actually show that e(ξ) is a well-defined function from 2-chains to **Z**, up to the choice of a 2-coboundary, i.e., the difference of functions defined by two choices of sections over 1-simplices is a 2-coboundary. To show that e(ξ) 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(ξ)$\in$.

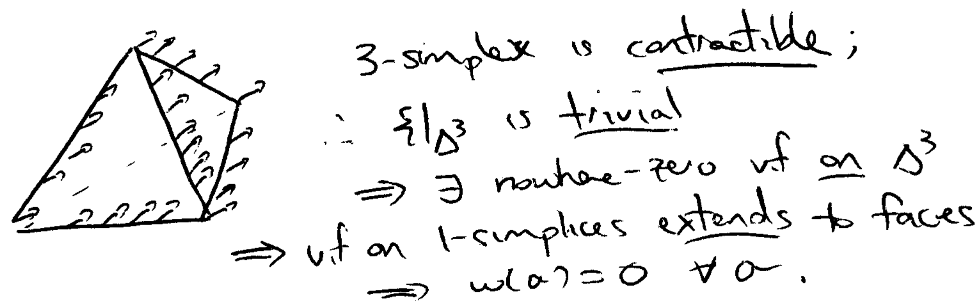

From this definition it is easy to see that e(ξ) is natural; if F is a subcomplex of M (e.g., an embedded surface in M), then $e\xi|_F = i^*(e(\xi) \in H^2(F)$, where i^* is the map on cohomology induced from the inclusion map i. This is because the winding number around the boundary of a 2-simplex will be the same, 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) surface F can be calculated by counting indices of a singular vector field on F. To see this, just remember that $H_2(F)=\mathbf{Z}$, generated by the fundamental class [F] (which is just the sum of all of the simplices of F, for some triangulation), so $e(\xi)\in H^2(F)=\mathbf{Z}$ 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 $S^1\rightarrow S^1$ which extends to $D_2\rightarrow S^1$, 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 ξ=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 comes 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: $[S]=[F]\in H_2(M)$. 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 $e(T\mathcal{F}|_F)[F] = e(TF)[F] = \chi(F)=2\text{-}2\text{genus}(F)$. 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

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

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 single 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 transverse 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, both normals are orthogonal to the vector orthogonal to the leaf, so are both tangent 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 $i^*(\mathcal{F})$.

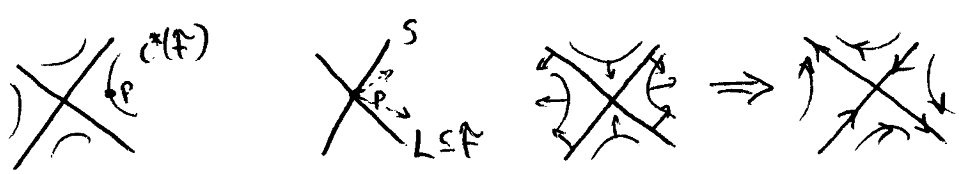

Notice that this vector field is tangent to S (by construction), but because it is tangent to the leaves of the singular foliation, it can also be thought of as a vector field in $\mathbf{T(\mathcal{F}|_S)}$. This is basically because locally it looks like the picture below. The point is that this one 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**.

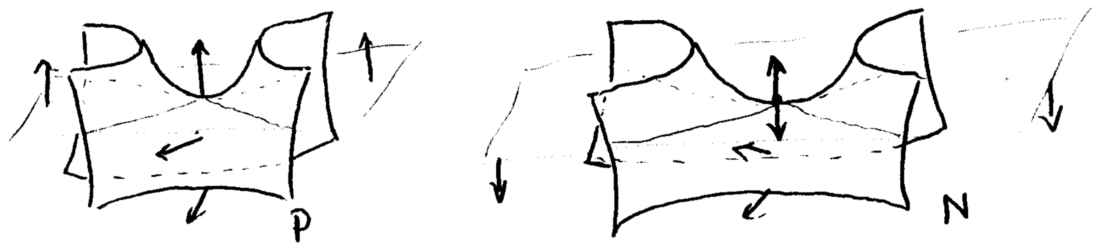

Let I_P = the number of positive saddles of $i^*(\mathcal{F})$, and let I_N = the number of negative saddles.

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!

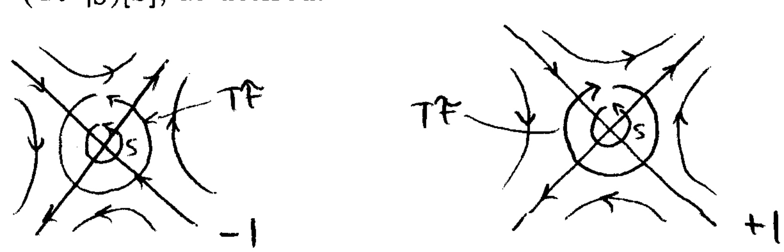

This finishes the case that $i^*(\mathcal{F})$ has no circle tangencies. If there are circle tangencies, we will show that we can still build <u>two</u> vector fields, one in T(S), the other in $T\mathcal{F}|_S$, which each agree with the one described above outside of small neighborhoods of the circle 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.

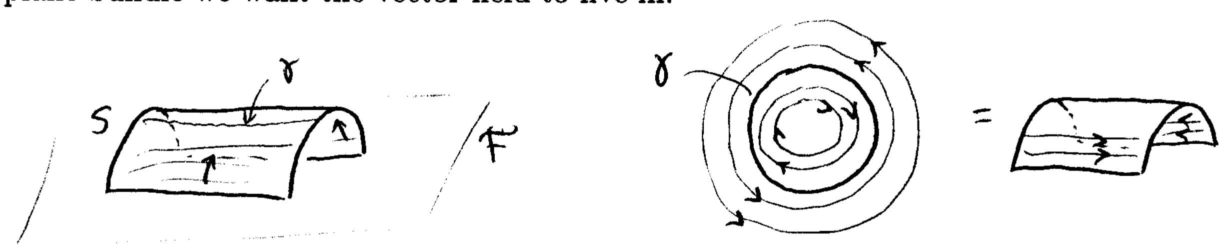

If we wish to create a vector field in T(S), this is relatively easy. We just alter the vector field between two loops on either side of γ as in the figure below. Basically, we are replacing the singular foliation $i^*(\mathcal{F})$ in between the loops to a 'Reeb annulus' like the one below, and taking 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 $i^*(\mathcal{F})$.

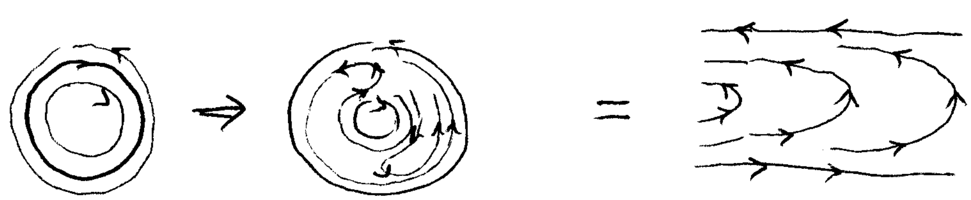

To build the vector field in $T\mathcal{F}|_S$, we need only be slightly more clever. The local picture is as below, with the vector field as above away from the circle tangency γ. This time we want to extend the vector field over the same little annulus around γ, except this time we want to make sure our vectors 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 horizontal vectors which go from pointing left to pointing right by 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.

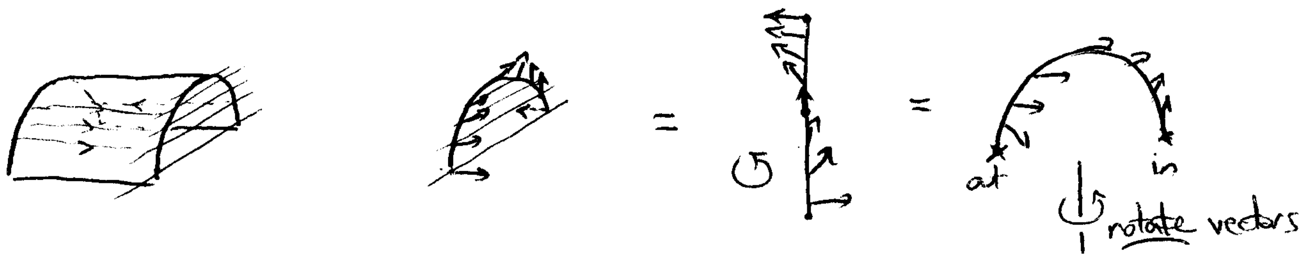

We finish with a few comments about this theorem:

(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 occuring 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 $\rightarrow$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 $F:[X,A]\rightarrow H^n(X;G)$, 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 $f^*:H^n(A;G)\rightarrow H^n(X;G)$ (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=**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(**Z**,1) is easy to find, namely S^1; $\pi_1(S^1)=\mathbf{Z}$, and since its universal cover is ($\mathbf{R}^1$ hence) contractible, all higher homotopy groups (agree with those of $\mathbf{R}^1$ 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,∂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,∂M),(I,∂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 $x:H_2(M)\to\mathbf{Z}$ 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 surger 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.

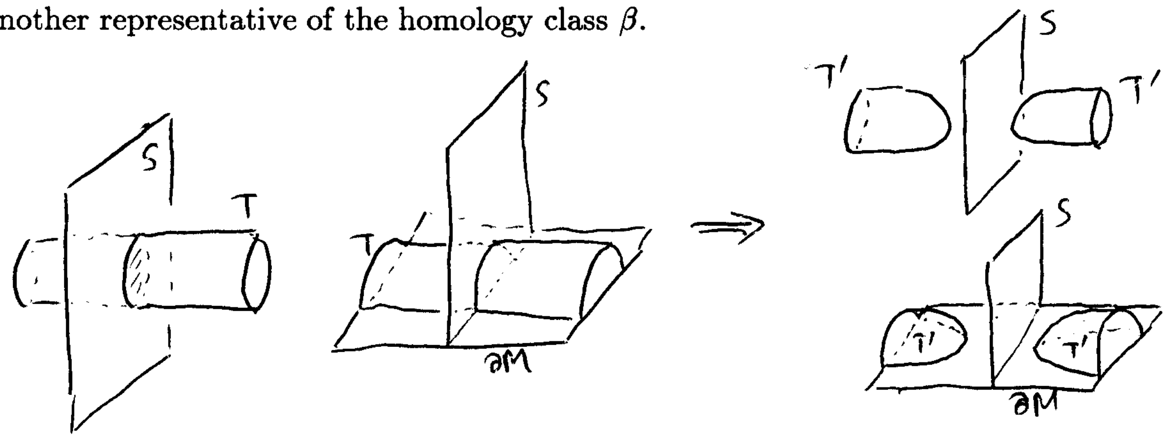

$\chi(T') = \chi(T) + 2$ or 1 (depending one whether we surgered along a circle or arc - this is easy to check), so we would expect χ_- to have gone down. This can't happen ($\chi_-(T)$ is minimal, so it must be that we have created new sphere or disk components by this process. But then it is easy to 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 homeomorphic 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 not 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∩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∪T)\(S∩T) (see figure); but one of those pieces must be an outermost disk or innermost disk, contradicting our preliminary surgery process. So the <u>same</u> collection of spheres and disks are thrown away on each side of the equality (*), giving an equality of χ_-'s, and finishing our proof.

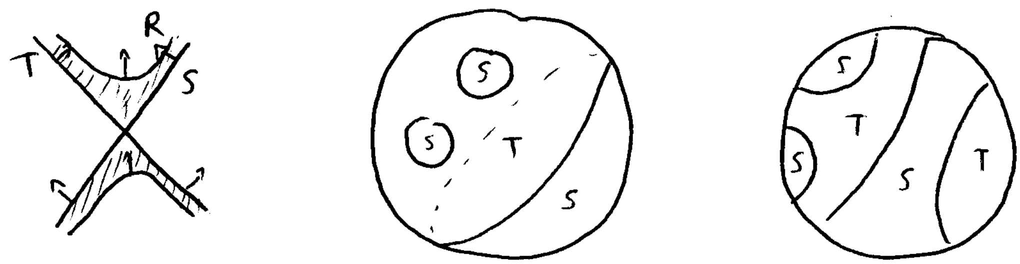

Next time we will handle part (b) of the proposition, and continue to explore some of the properties of this ~~metric~~. The next step is to show that these norms extend, in an essentially unique way/ to continuous norms for $H_2(M;\mathbf{R})$ and $H_2(M,\partial M;\mathbf{R})$, taking values in $\mathbf{R}$.

norm! Wrong class.

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 $n\alpha$ 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 $M\setminus N(T)$ 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 $g:M\to S^1$, and if we then compose g with the standard n-fold covering map $p:S^1\to S^1$, we also get a map $p\circ g:M\to S^1$ representing $n\alpha$. 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 pog are homotopic to one another. But since p is a covering map, the Covering Homotopy Theorem implies that there is a map g':M$\to S^1$ 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(α)=0 implies α=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 χ_-=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:$H_2(M;\mathbf{R})\to\mathbf{R}^+$ (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(kα) is defined. If we are to define x(α) 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 <u>define</u> x(α); 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(kα), so the definition of x(α) 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 kα and kβ 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$\mathbf{Q}^n\to\mathbf{Q}^+$ 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, $x(a\alpha+b\beta)=0$. 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 <u>unit</u> <u>ball</u> 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 ξ 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 <u>polyhedra</u>, i.e., they are each the intersection of a finite number of half-spaces.

This is actually true for <u>any</u> 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 $k=1$ 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

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

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

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

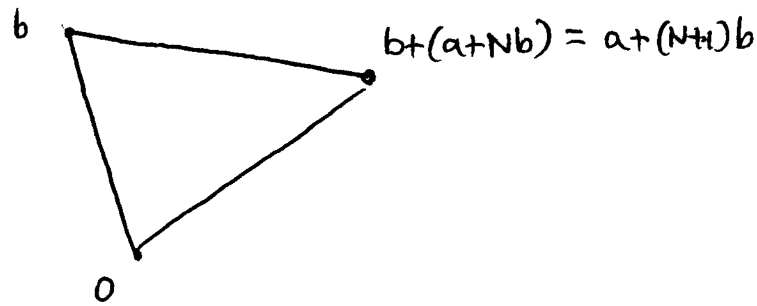

But then because x is linear on rays through the origin, it is easy to see that x is therefore linear on the cone on this line segment [b,a+(N+1)b] in $\mathbf{R}^n$. This is the basic building block of our construction. What we will do is show that, for any indivisible element b of $\mathbf{Z}^n$ (meaning it can be extended to a **Z**- basis $b=a_1,a_2,\ldots,a_n$ of $\mathbf{Z}^n$) x is linear on the convex span of some (**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 **Z**-basis of $\mathbf{Z}^{k+1}$) 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 $\mathbf{R}^n$ 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 $\mathbf{Z}^n$ (and hence of $\mathbf{R}^n$) so that x agrees with a linear function L_b on their convex span $C=C_b$.

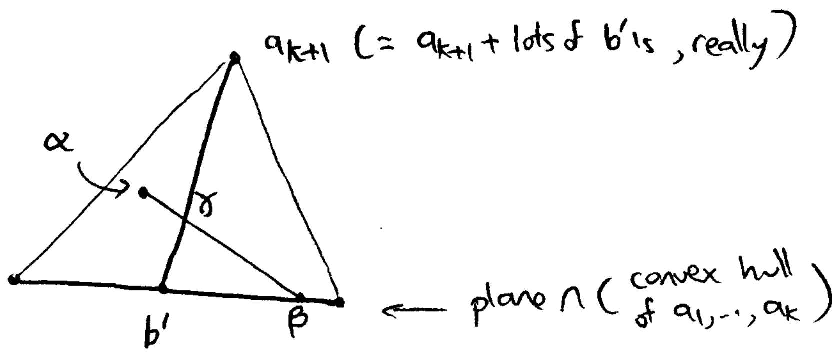

Notice that we also get that x agrees with $-L_b$ on the convex span of $-a_1,\ldots,-a_n$; both functions change sign. But now since x and L_b agree on this cone, they are both equal to 1 at the same points of this cone. For L_b, these points consist of the intersection of an affine hyperplane P_b with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

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 $\alpha\in\mathbf{R}^n$, $x(\alpha)\leq 1$, 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 <u>interior</u> of the cone where x and L_b agree (but they don't agree!).

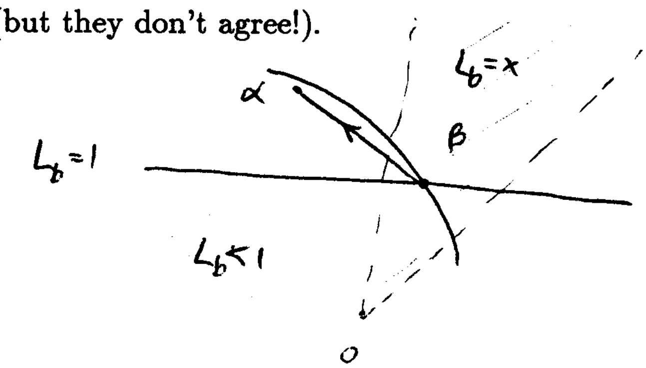

Note that this also says that if $x(-\alpha)=x(\alpha)\leq 1$, then $L_b(-\alpha)=-L_b(\alpha)\leq 1$, i.e., $x(\alpha)\leq 1$ implies $|L_b(\alpha)|\leq 1$. Consequently, if we think of L_b as an element of the dual space $\mathrm{Hom}(H_2(M,\partial M;\mathbf{R}),\mathbf{R})$, it has $x^*(L_b)\leq 1$. 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 $x^*(L)\leq\Sigma_{i=1}^k t_i x(L_{bi})=\Sigma t_i=1$ (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.