Foliations and the Topology of 3-manifolds

Outline of class 21

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

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

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

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

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

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

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

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

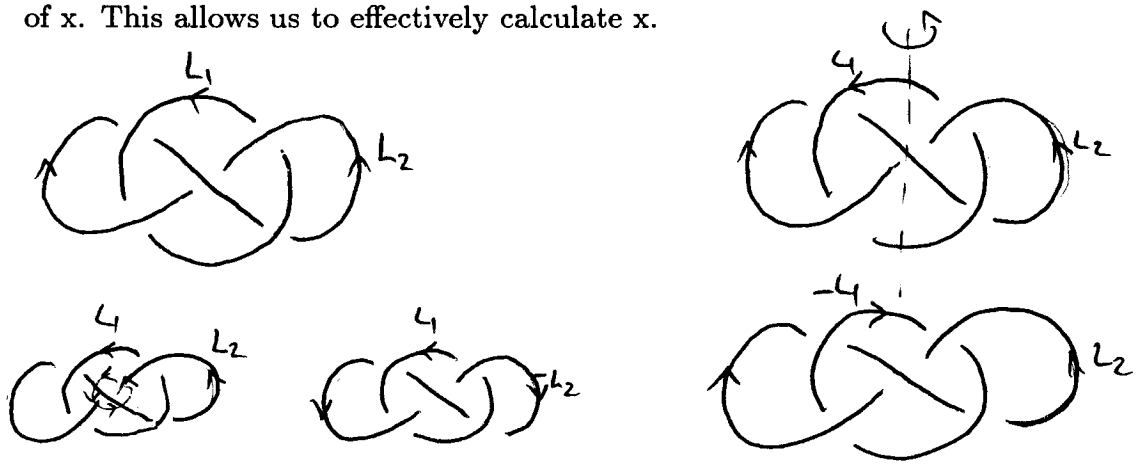

Start with $x(L_1)$ $(=x(L_2))$. We need to find a surface whose boundary represents L_1 in $H_1(N(L))$. Inspection quickly turns up one of the surfaces S,T below; the first is a once-punctured torus, and the second is a twice-punctured disk. The second is really no better than the first - it's obtained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

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

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

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

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

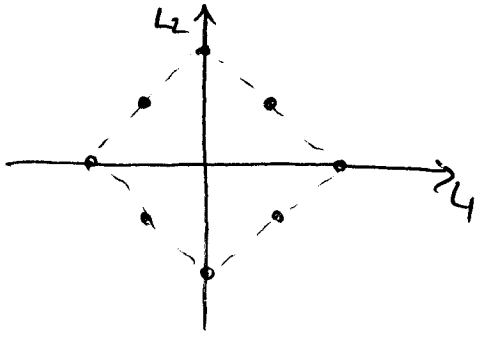

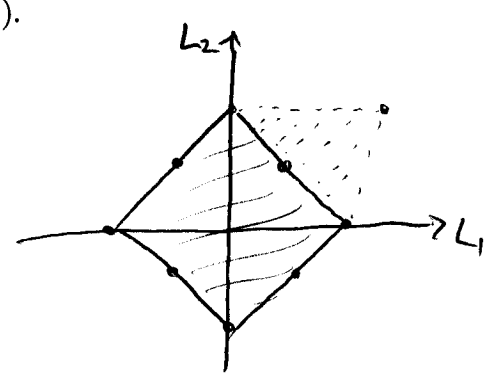

Notice that we in fact proved that B_x as a polyhedron, in this case, instead of using it to help us find out what it looked like. This is something that can be used in general. If you already have (somehow) a collection of points which you know are in ∂B_x, and they include both a collection (the vertices) of points P and a point in the interior of each of

the top-dimensional faces of the convex hull of P, then since the only convex set which contains all of these points in its boundary IS the convex hull of P, this must be exactly the unit ball of the norm.

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

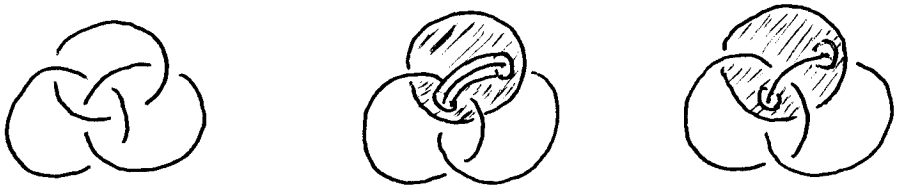

Again, the Borromean rings has a great deal of symmetry, which we can exploit in our calculations. If we rotate the picture around its center, we can permute the components $(1,2,3)\mapsto(2,3,1)$ and to $(3,1,2)$. By rotating around lines in the plane which are vertical, and inclined 120 degrees, we get the permutations $(1,2,3)\mapsto(-1,-3,-2)$, $(-3,-2,-1)$, and $(-2,-1,-3)$. Finally, by using a different view of L below and rotating about the z-axis, we get $(1,2,3)\mapsto(1,-2,+3)$ (& similarly for other axes), while by pushing L to this position and over turns each component inside out, so we then get $(1,2,3)\mapsto(-1,-2,-3)$, and then composing with another one above gets us to $(1,3,2)$. So we can achieve any rotation or transposition or change of sign, hence our norm will have the corresponding symmetries.

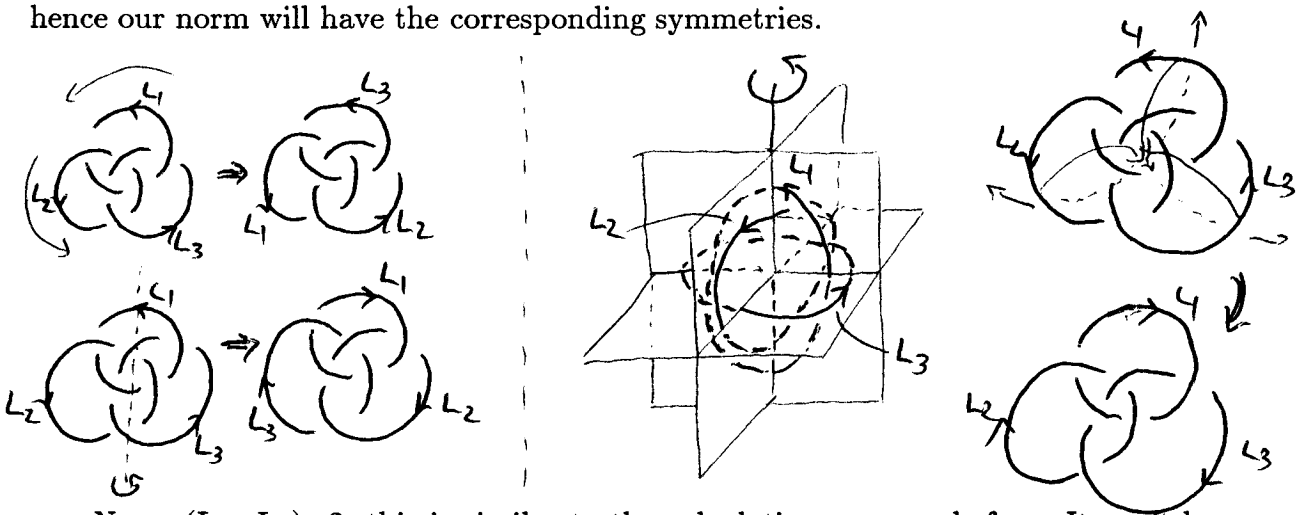

Now $x(L_1+L_2)=2$; this is similar to the calculation we gave before. It can't be zero since then symmetry implies $2x(L_1)\leq x(L_1+lt)+x(L_1-L_2)=0$. Also, it can't be 1; since a χ_--minimzing surface S has at least two ∂-components (and an even number of components altogether on $\partial N(L_1)\cap\partial N(L_2)$), and an odd number of ∂-components (since $\chi(S)$ is odd),

there must be a meridian on L_3; by capping off with a disk, we get a surface bounded by $L_1 \cup L_2$ intersecting L_3 exactly once, again contradicting linking number 0.

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

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

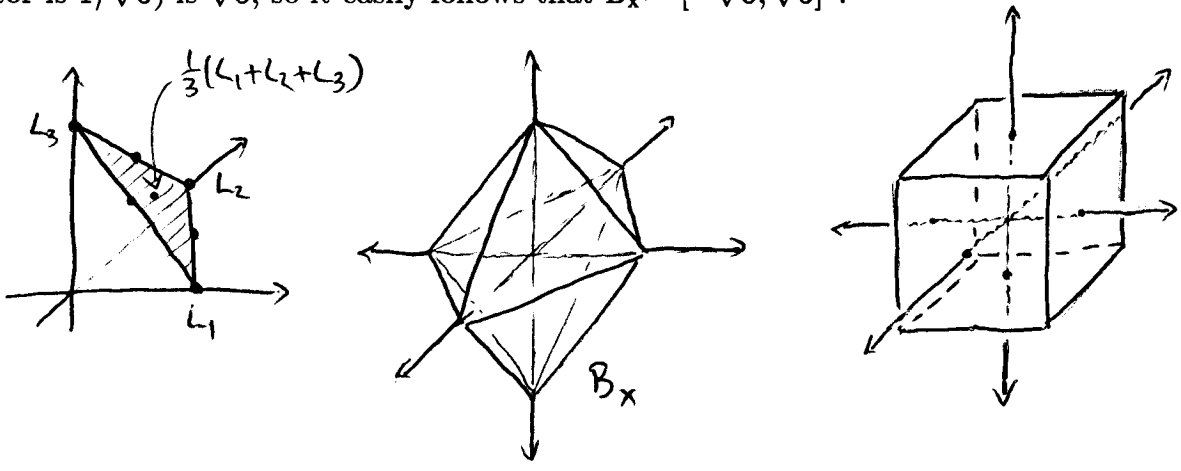

Next time we will do a somewhat more complicated example (another 3-component link), and then start looking at the mechanics behind Gabai's construction of foliations using the Thurston norm.

Foliations and the Topology of 3-manifolds

Outline of class 22

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

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

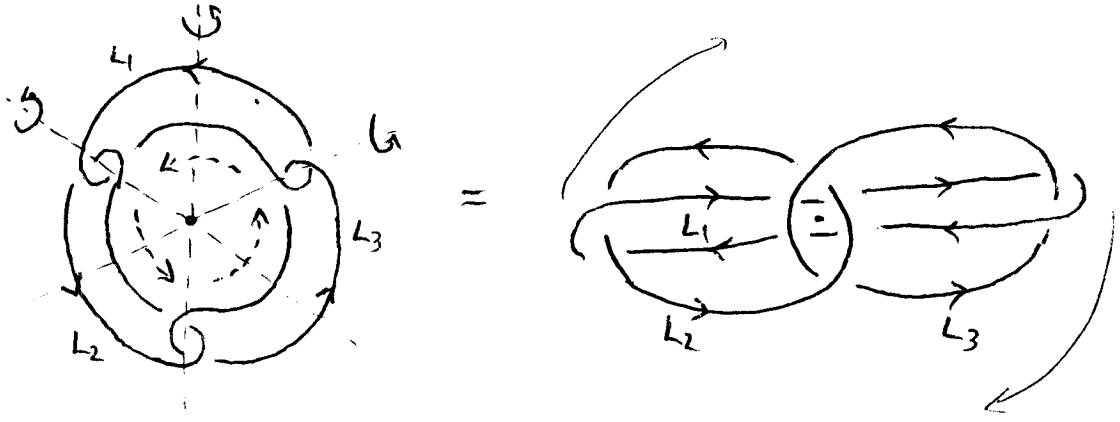

Now let's calculate. As before, $x(L_1)=1$; a surface with this χ_- is given below, so if $x(L_1)\neq 1$, then $x(L_1)=0$, so L_1 could be represented by a disk or annulus. But, as before, a disk contradicts the non-triviality of the link (or, if you prefer, the fact that the linking number of pairs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

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

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

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

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

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

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

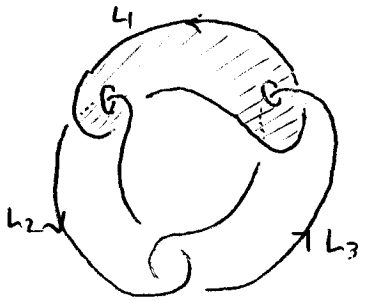

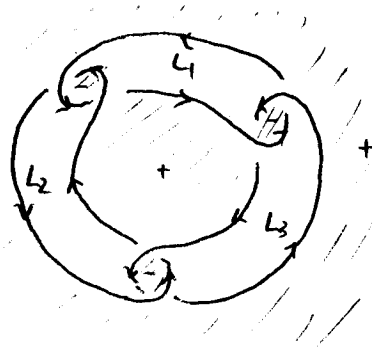

Therefore $x(\pm L_1 \pm L_2 \pm L_3)=1$ if all of the signs are the same, otherwise it is $=3$; and $x(L_i \pm L_j)=2$ for all $i \neq j$. If we plot all of the points this then guarantees have norm 1 (see below), we find that their convex hull is basically the octagon that formed the unit ball for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

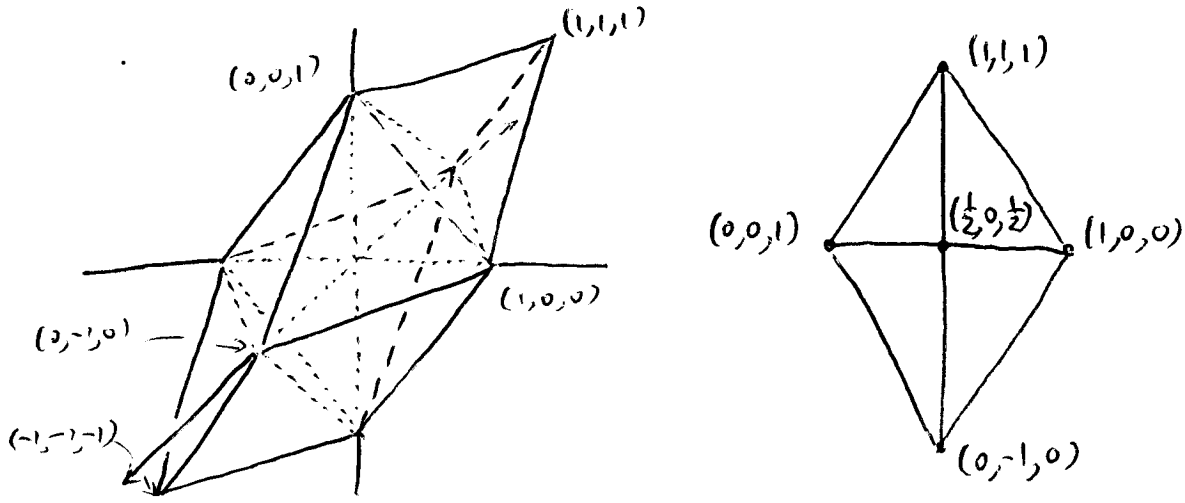

So that's one final, somewhat more interesting example.

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

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

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

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

Definition: A **sutured manifold** is a pair (M,γ), where M is a compact orientable 3-manifold, and $\gamma\subseteq\partial M$ is a collection of disjoint annuli and tori, denoted $A(\gamma)$ and $T(\gamma)$. Each annulus A comes with an orientation, in the form of an orientation to its core circle, called the suture, and denoted s(A); the collection of sutures is denoted $s(\gamma)$. Each component of the rest of the boundary, $R(\gamma)=\partial M\backslash \text{int}(\gamma)$, can be given a (normal) orientation, allowing us to split $R(\gamma)$ into two pieces; $R_+(\gamma)$ consist of those pieces for which the normal points out of M, and $R_-(\gamma)$ consists of the pieces with inward-pointing normal. These

normals, however, must be chosen in a way consistent with the orientations of the sutures - if we orient a component of $R_{\pm}(\gamma)$ (using the normal orientation as third vector of an orientation for M), this induces an orientation on its boundary components; these must agree with the orientations on the suture in the bordering annulus. See the figure below.

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

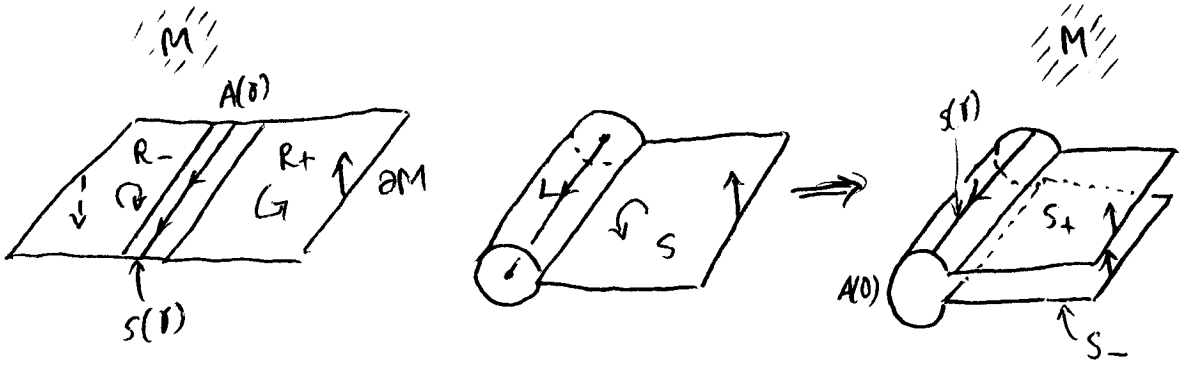

One thing to notice: the orientation conventions of a sutured manifold require that as you cross a suture you must pass from R_+ to R_- (or vice-versa) - the orientation on the boundary reverses itself as you cross a suture (from the point of view of ∂M).

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

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

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

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

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

Foliations and the Topology of 3-manifolds

Outline of class 23

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

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

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

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

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

Figure 1

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

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

Figure 2 Figure 3

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

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

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

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

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

Figure 4

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

As an example, take the (Hopf) link L pictured in Figure 5 (which we think of as sitting in S^3). We can easily find a spanning surface for L, and if we imagine thickening up this (annular) surface, we get a solid torus (whose complement, which is actually our sutured manifold, is also a solid torus). We can then draw the (cores of the) resulting annular sutures on this solid torus, to gives us an inside-out view of our sutured manifold. Remembering that the disks we want to find should lie outside of the solid torus we see in from of us, it's not hard to find a candidate for a decomposing disk. All that remains is to figure out what it looks like when we decompose along this disk. It's clear that the boundary of our new sutured manifold will be a sphere (and our sutured manifold will be a solid torus cut open along a meridian disk, i.e., a ball) - we must then figure out what

the new sutures look like. But this is actually not to hard to figure out. The boundary of the (oriented) disk comes with its own orientation, as do the two original sutures, and the new sutures inherit these same orientations, so there is really only one way to glue the pieces together in a way that respects orientations - see the figure. Inspection shows that the resulting suture is connected, so the sutured manifold we get is actually $D^2 \times I$. Therefore the sutured manifold we started with is disk decomposable.

Figure 5

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

Figure 6

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

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

Figure 7

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

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

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

Figure 8 (no joke intended)

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

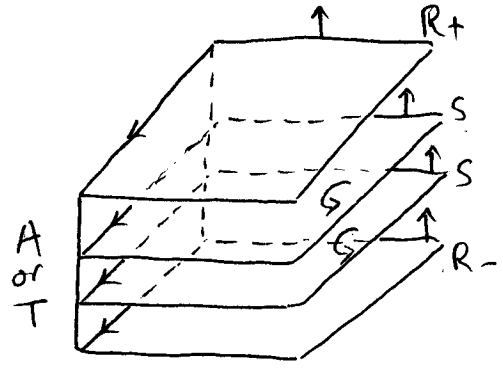

Fig. 1a

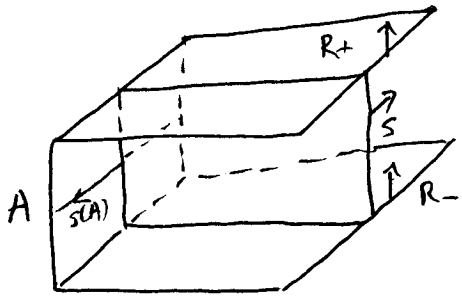

Fig. 1b

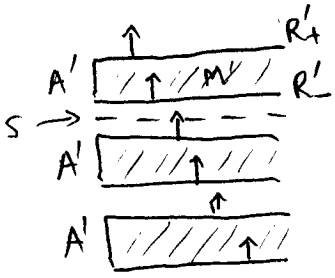

Fig. 2

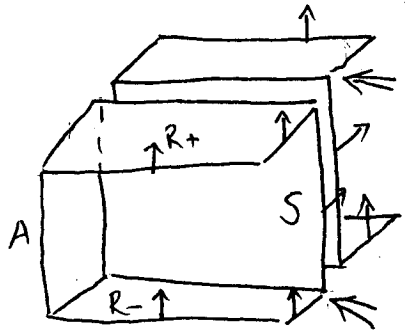

Fig. 3a

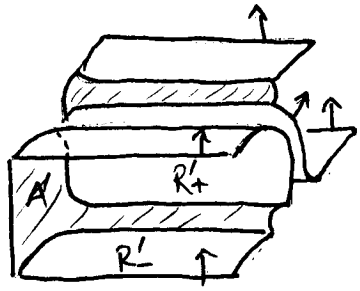

Fig 3b

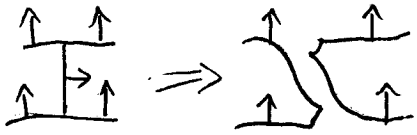

Fig. 3c

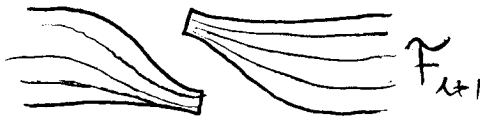

Fig. 4a

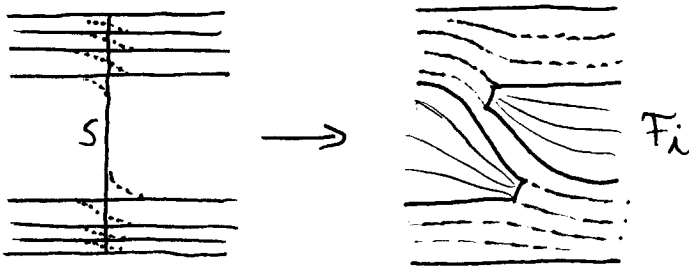

Fig. 4b

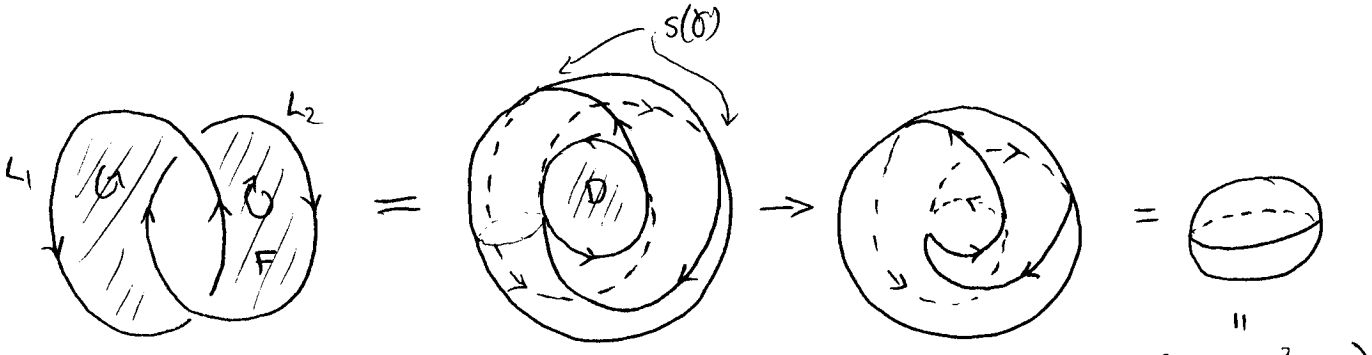

Figure 5

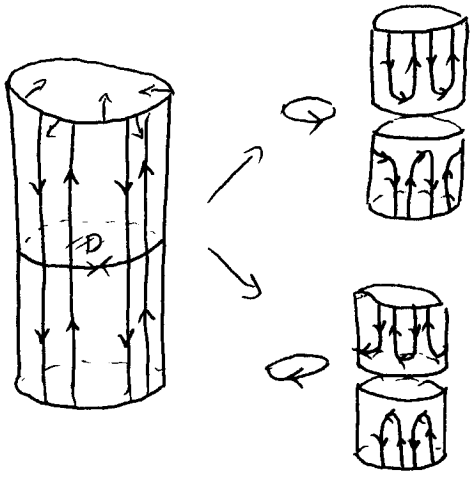

Figure 6

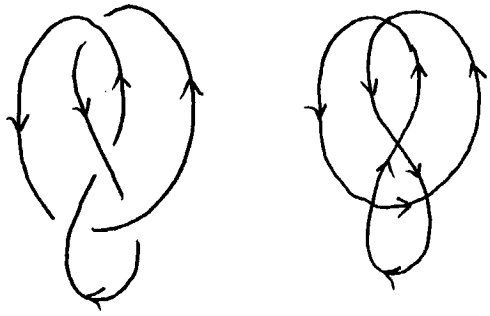

Fig. 7a

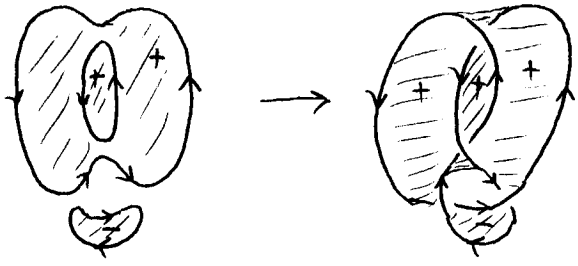

Fig. 7b

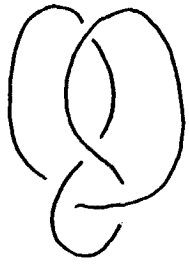

Fig. 8a

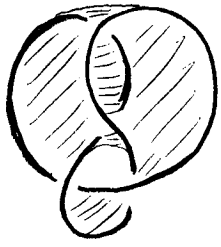

Fig. 8b

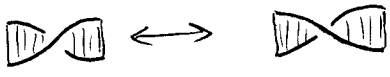

Fig. 8c

Foliations and the Topology of 3-manifolds

Outline of class 24

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

Our goal is to prove:

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

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

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

Figure 1

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

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

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

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

In the first case, a hanging edge means a Seifert disk with only on band attached to it; by flipping the disk over, we can amalgamate two disks and a band into a single disk. Clearly, the projection remains alternating. In the second case, there is a loop in the plane which hits Γ once in that edge. this loop separates the plane into two disks, and by flipping the entire link projection in one of them over, we again amalgamate two disks

and an edge (the offending one) into a single disk. Again it is not hard to see the resulting projection of our link is still alternating, and further that the resulting surface is the same as we would get by running Seifert's algorithm again on the new projection (Figure 6). Since these actions in both cases reduce the number of crossings in our projections, we eventually can't do either. (Those of you attending Cameron's class no doubt recognize this as the process of getting a reduced diagram for an alternating link.)

Figures 3 through 6

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

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

Figures 7 through 11

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

Figure 12

That completes the proof in the first case (actually, it really does! The resulting link also has non-nested Seifert circles - they're basically the same ones we started with (modulo

the amalgamation).). To deal with the second case (the circles are nested), we need to introduce a new concept.

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

Figure 13

If we (normally) orient both F_1 and F_2 to point upwards on their common disk D, then together they give a normal orientation for F. This in turn gives a sutured manifold structure to $M=S^3\backslash \text{int } N(F)$. We will prove our theorem in our second case by showing that the Seifert surface is the Murasugi sum of two surfaces, each of which is disk decomposable. Then the following result will finish our proof:

Proposition: If F_1 and F_2 are both disk-decomposable, then so is F.

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

All we need is to find that one decomposing disk. But if we think of our picture of S^3 above as being $S^2\times R\cup\pm\infty$, then there is an obvious disk to try - namely the complementary disk D′ to the disk D in the equatorial sphere (see Figure 13). All we have to do is see that it works! But we have (by having an orientation on the two pieces F_1 and F_2) already given an orientation to the boundar of (D, hence) this disk D′. Even though we have the option of choosing one of two orientations for this disk, we would be fools to simply ignore the one that the setup hands us, so we might as well see what happens if we decompose using that orientation. This is shown in Figure 14; basically, the two pieces lift off of one another. It only remains to figure out which way the two pairs of vertical sutures are joined together by the mew pieces of suture that come from D′; but it is easy to see that they basically just give us our original ∂F_1 and ∂F_2 back again. Finally, by isotoping D′ in each piece around (not at the same time! Pretend we have only one piece embedded in S^3 at a time) to lie parallel to D, we easily see that each piece is homeomorphic to $S^3\backslash$(the appropriate thickened surface), as desired (Figure 15).

Figure 14 Figure 15

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

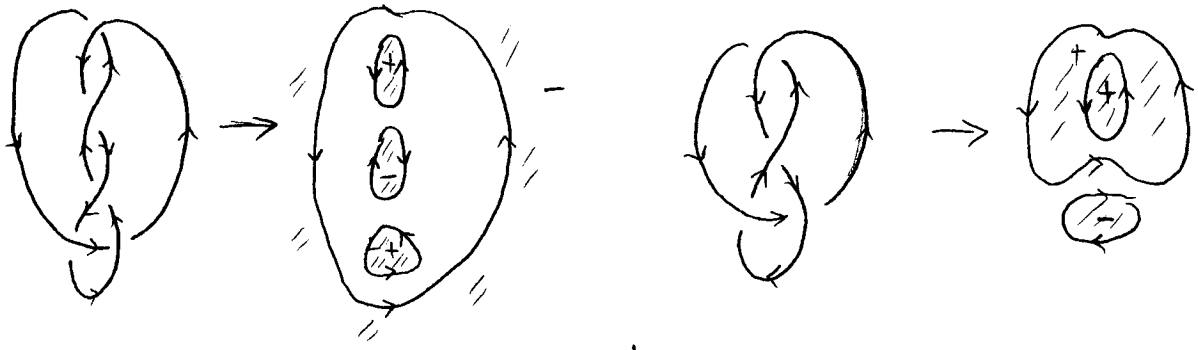

Figure 1

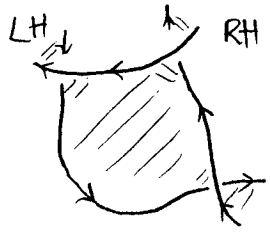

Fig. 2

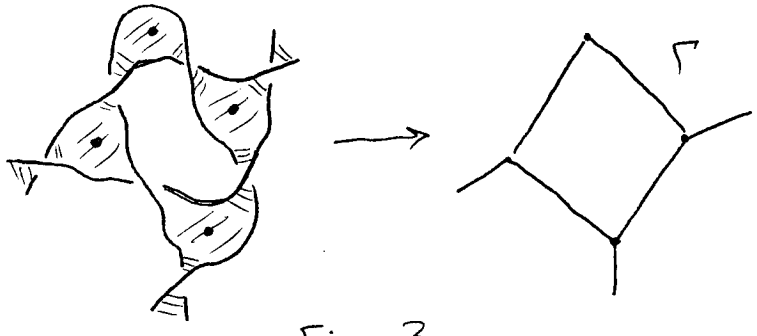

Fig. 3

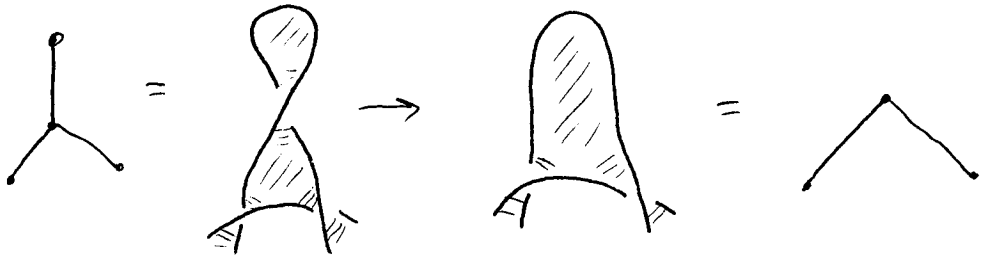

Fig. 4

Fig. 5

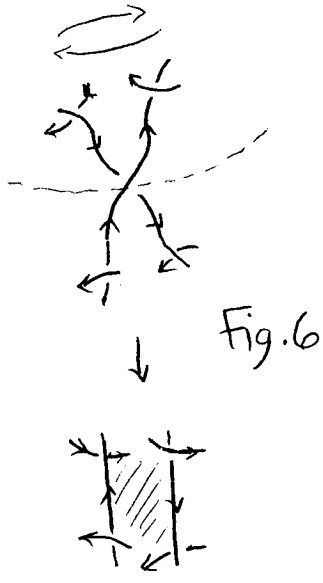

Fig. 6

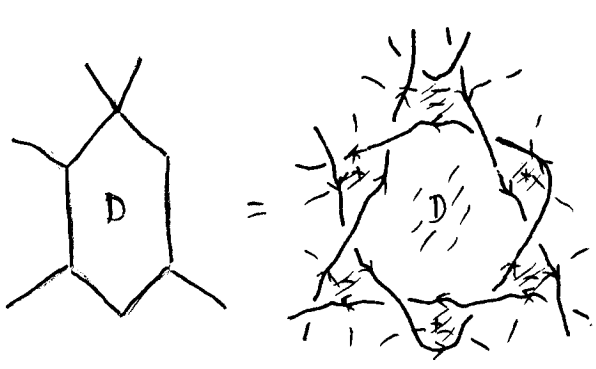

Fig. 7

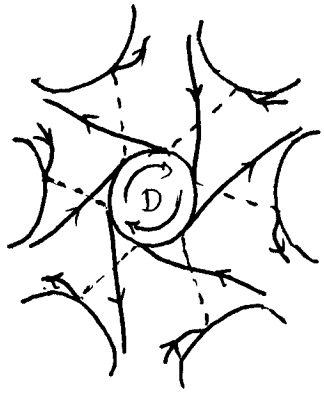

Fig. 8

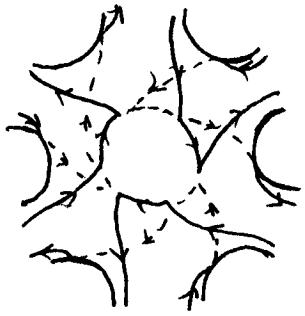

Fig. 9

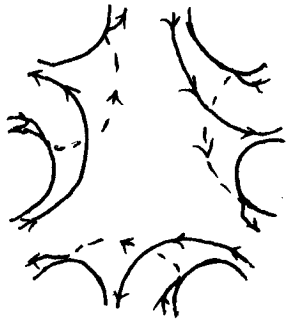

Fig. 10

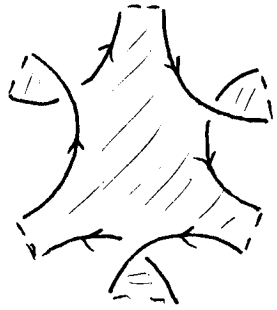

Fig. 11

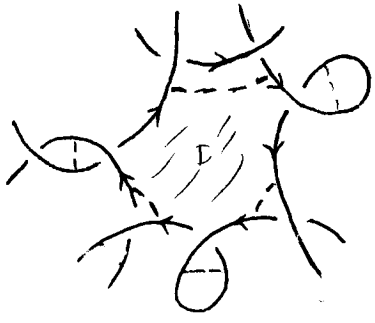

Fig. 12

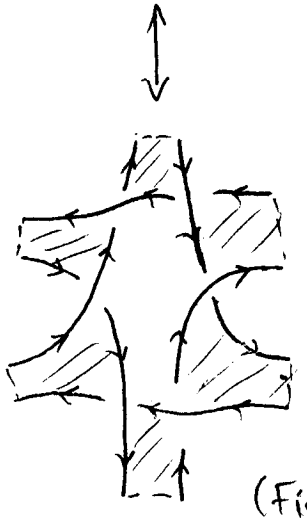

(Fig. 7)

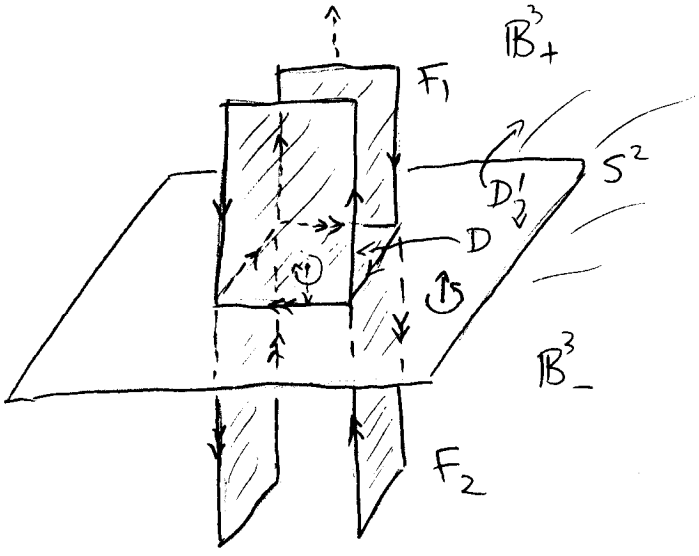

Figure 13

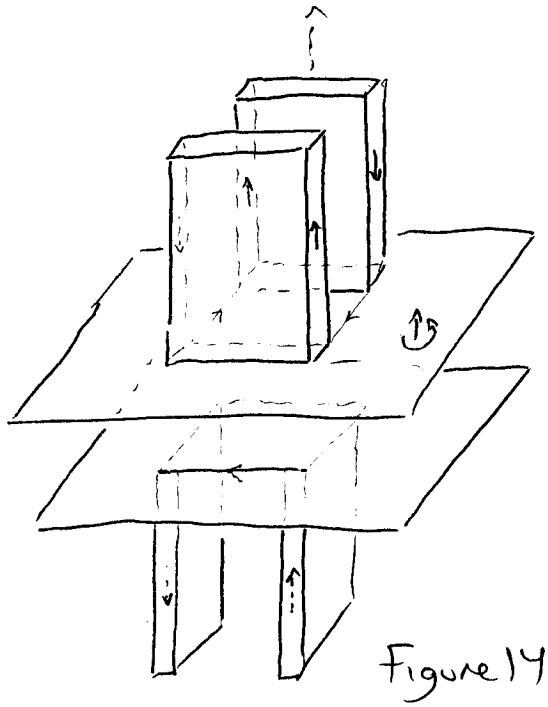

Figure 14

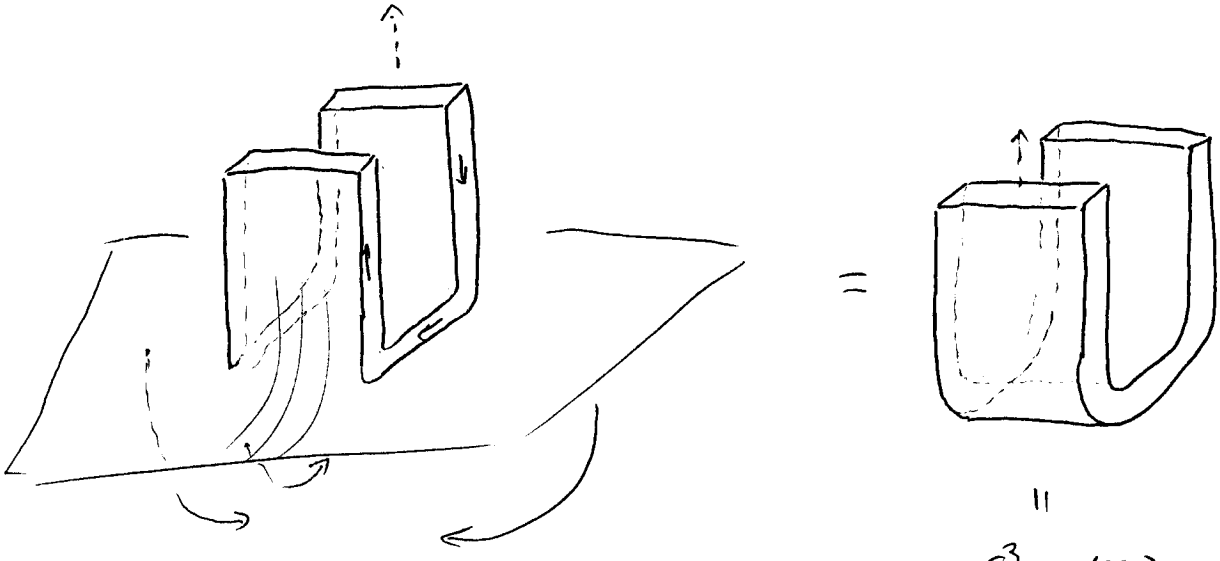

Figure 15

Foliations and the Topology of 3-manifolds

Outline of class 25

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

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

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

Figure 1 Figure 2

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

Figure 3 through Figure 5

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

Definition: A **pretzel link** P=P(n_1,...n_k) is a link obtained by stringing k twisted bands (each given n_i half-twists) together as in Figure 6. If the integer n_i is positive, we give positive twists; if negative, we give negative twists (Figure 7). It is not hard to see that P is not a knot if two or more of the n_i are even; this is because how many components the link has is determined not by the integers n_i but by their parity (Figure 8), so we

could imagine using 0 or 1-twisted bands, to count components, which makes the assertion obvious. If all n_i are odd, then P is a knot if k is odd, otherwise it has 2 components. Notice also that the ordering $n_1,\ldots,n_k$ is in part illusory - it should actually be thought of as a cyclic ordering, since we could instead draw our standard picture around a circle (Figure 8).

Figure 6 through Figure 8

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

Figure 9 Figure 10

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

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

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

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

Figure 11 through Figure 14

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

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

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

Figure 15

Case 2: All of the $|n_i| \geq 3$, and, for some i, $n_i \cdot n_{i+1} > 0$ (i.e., they both have the same sign).

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

Figure 16

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

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

Figure 17 through Figure 22

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

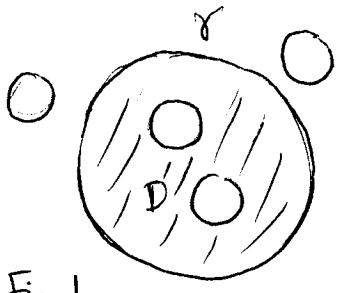

Fig. 1

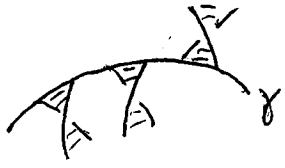

Fig. 2

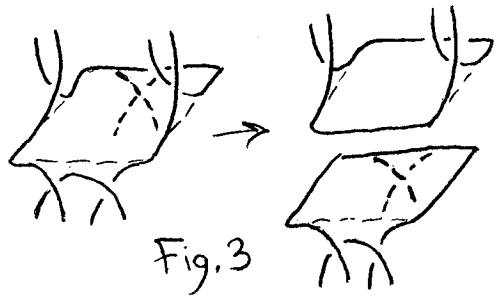

Fig. 3

Fig 4, 5

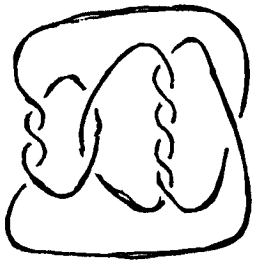

P(-3,3,-5,-1)

Fig. 6

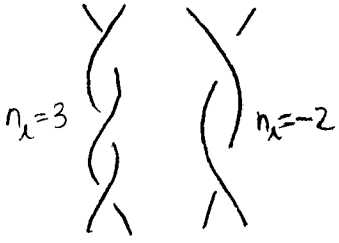

Fig. 7

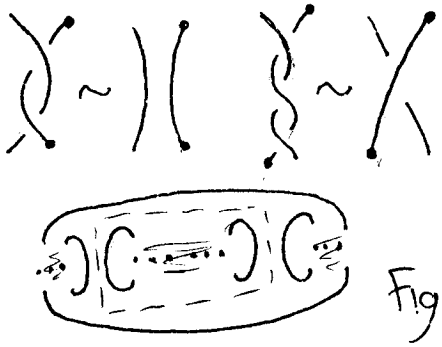

Fig. 8

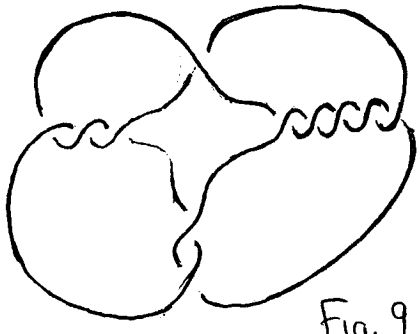

Fig. 9

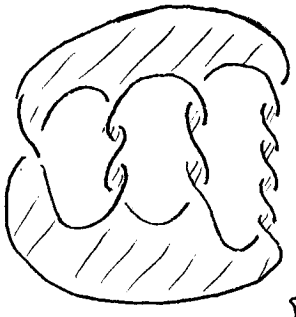

Fig. 10

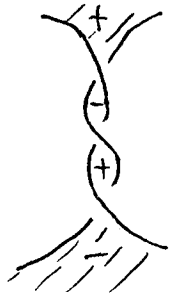

Fig. 11

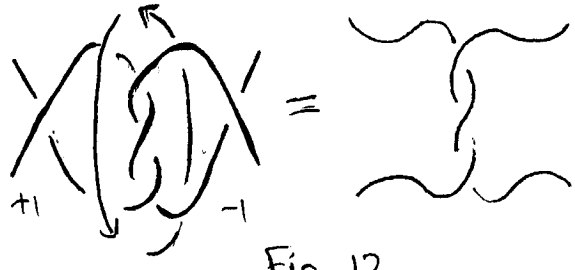

Fig. 12

Fig 13

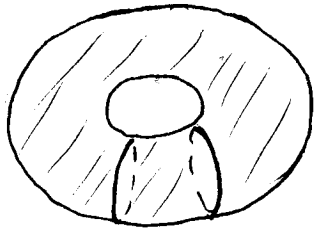

Fig. 14

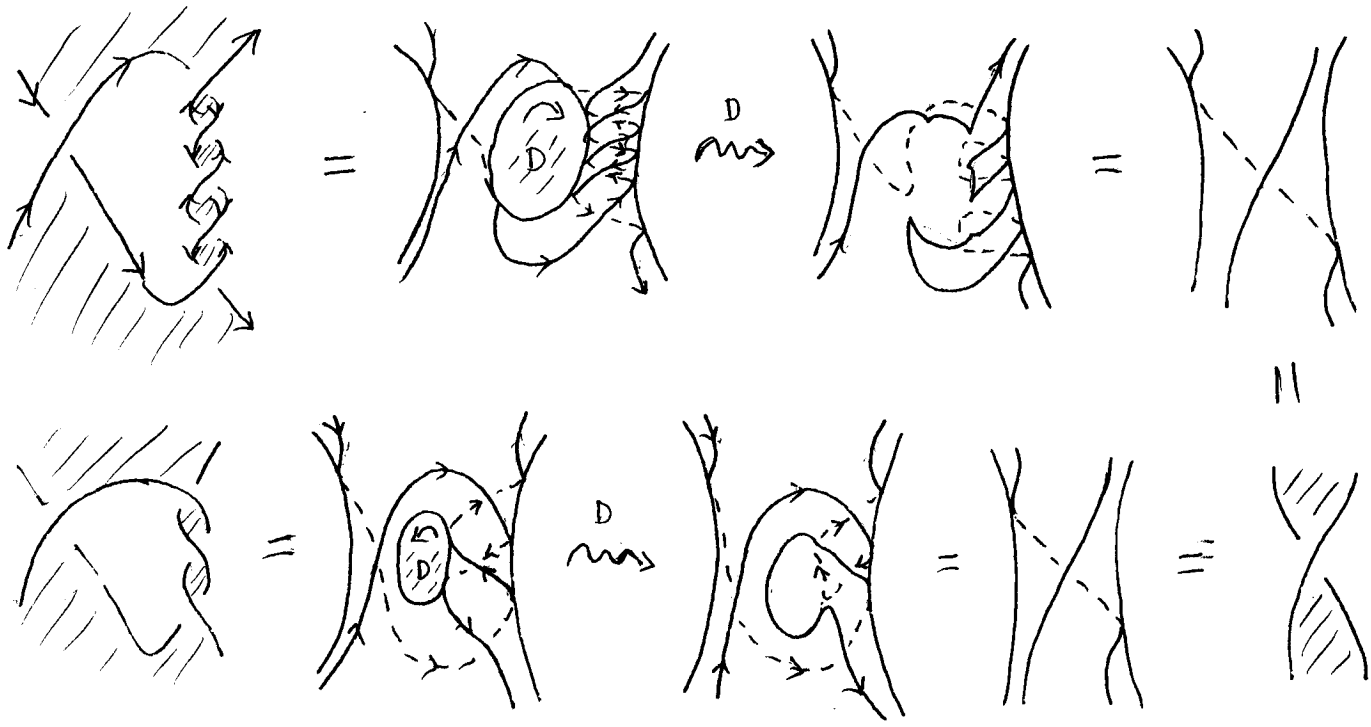

Figure 15

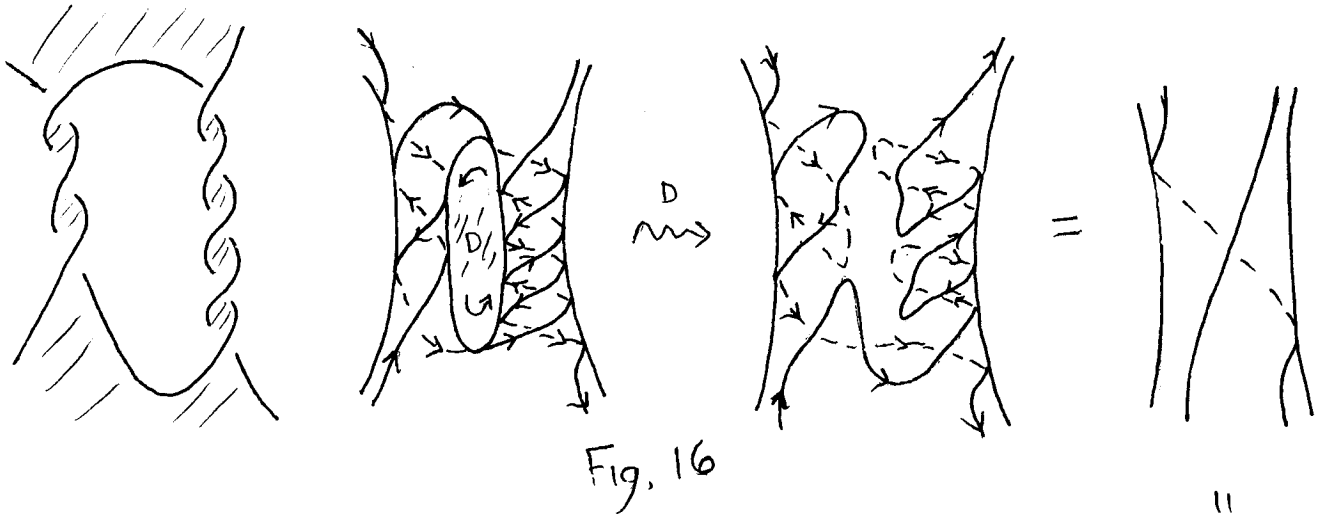

Fig. 16

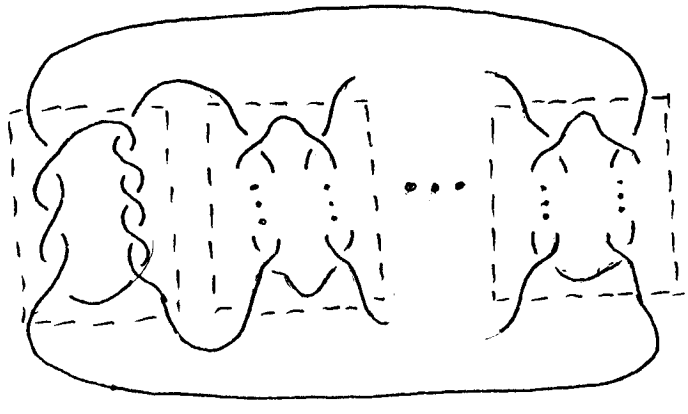

Fig. 17

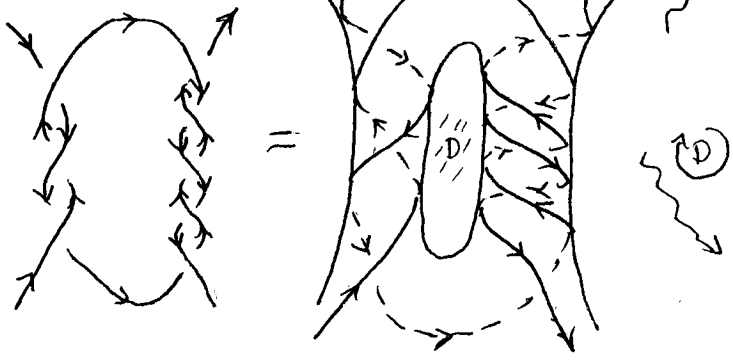

Fig. 18

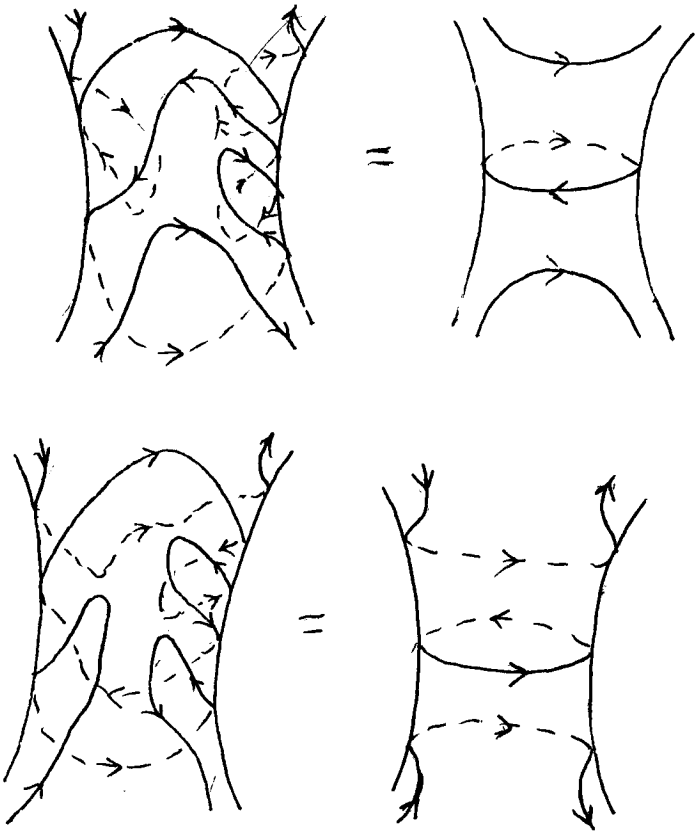

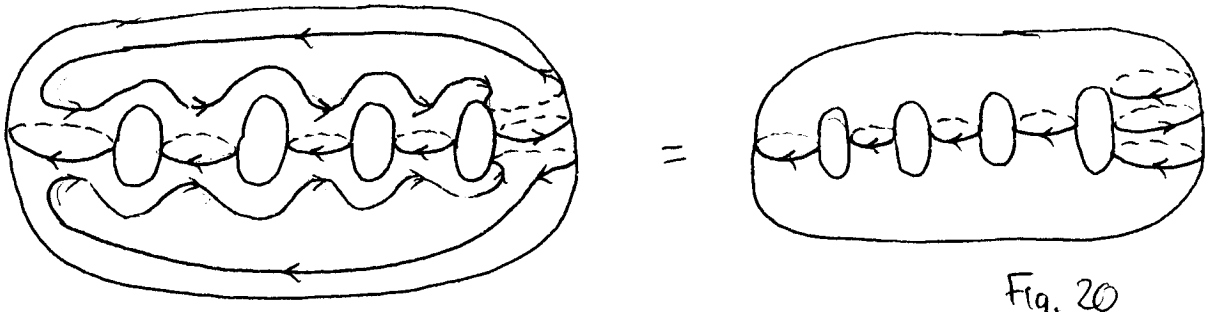

Fig. 19

Fig. 20

Fig. 21

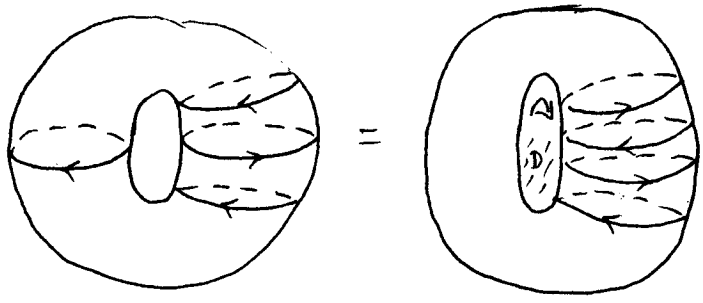

Fig. 22a

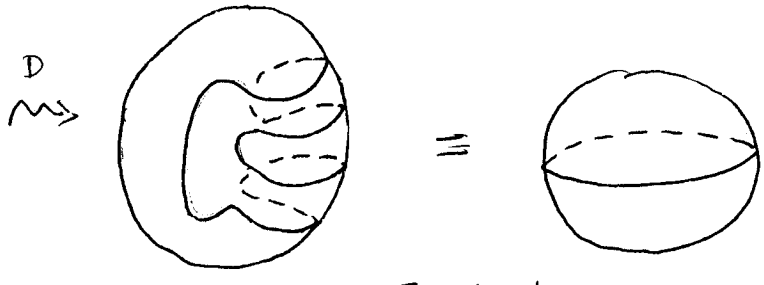

Fig. 22b

Foliations and the Topology of 3-manifolds

Outline of class 26

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

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

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

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

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

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

The corollary follows immediately from the two theorems and the definition of a disk decomposable surface. Now if we were to do slope 0 Dehn surgery (also known as 0-frame

surgery) on K, i.e., glue a solid torus to M so that the longitudinal circles bound meridional disks in the solid torus, then we get a new manifold M^+. By gluing meridional disks to the leaves of $\mathcal{F}$ along $\mathcal{F} \cap \partial M$, it is easy to see that we can build a foliation $\mathcal{F}^+$ of M^+, as well. But it is easy to see that this new foliation is taut, because $\mathcal{F}$ is: $\mathcal{F}^+$ has no new leaves, so the transverse loops that proved that /cafsp was taut suffice for /cafpl, as well. Notice that $F \cup D^2$ is a laef of $\mathcal{F}^+$.

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

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

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

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

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

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

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

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

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

We will do cases (2) and (4) first. In case (2), we have M_1 = a product 3-ball, so , since D intersects γ in 2 points, M is a product solid torus (Figure 8), i.e., an annulus×I. We can therefore foliate M by parallel annuli; but this would violate our extra assumption - this would not have leaf space S^1. But it is easy to see how to alter this product foliation to make it have leaf space S^1; just cut along our disk D and reglue in a way that preserves

the foliation by parallel lines, but shifts all of them 'up' (Figure 9). In case (4) we need not be so careful; D looks like a rectangle (Figure 10), so gluing N_1 to N_2 gives us something homeomorphic to N_2 (Figure 11); in fact (M,γ) is really just (N_2,γ_2), so we just use its foliation as the foliation of M.

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

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

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

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

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

Fig. 1

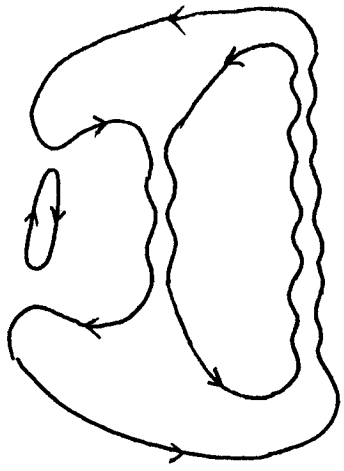

Fig. 2

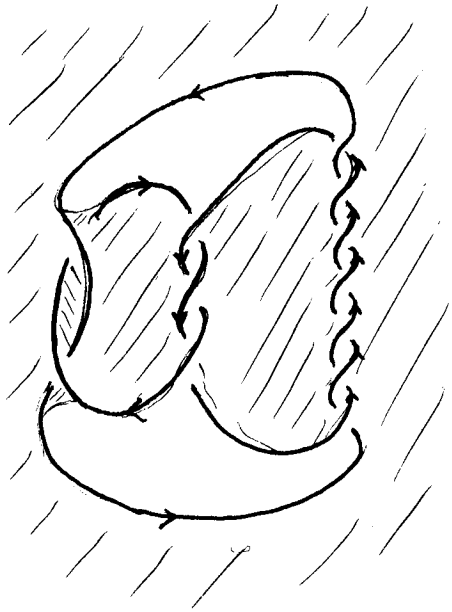

Fig. 3

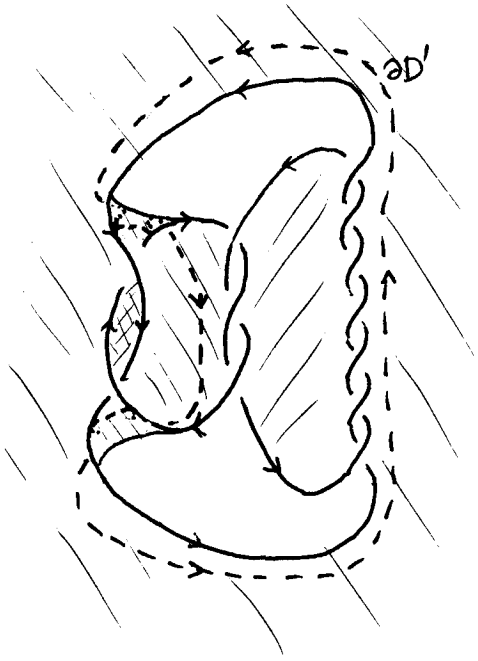

Fig. 4

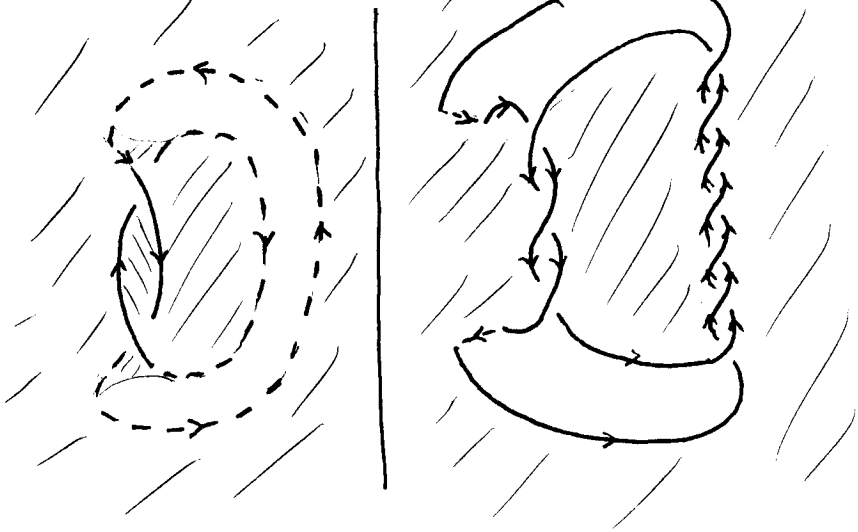

Figure 5

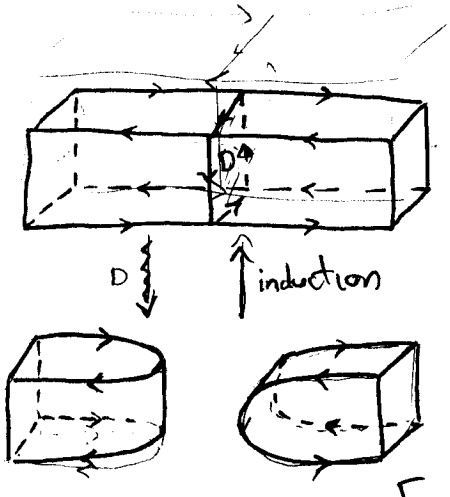

Fig. 6

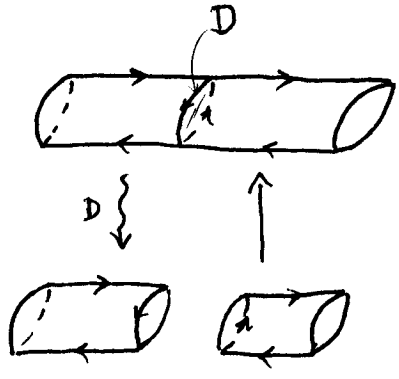

Fig. 7

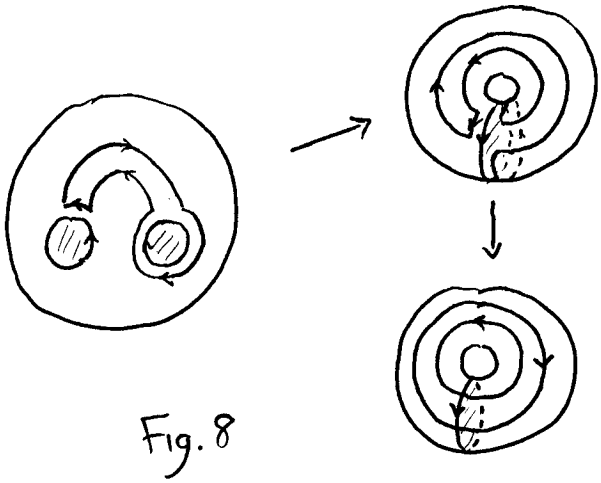

Fig. 8

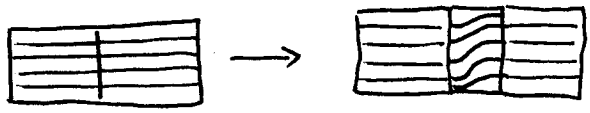

Fig. 9

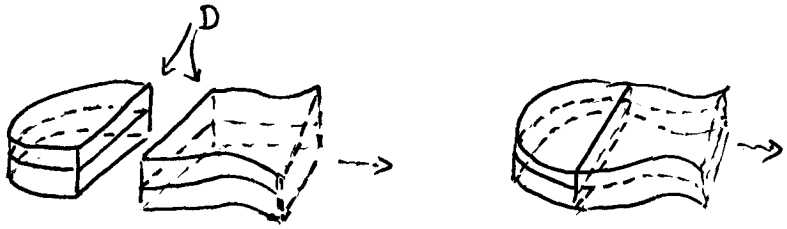

Fig. 10

Fig. 11

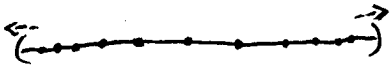

Fig. 12

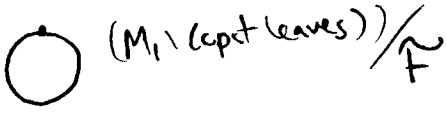

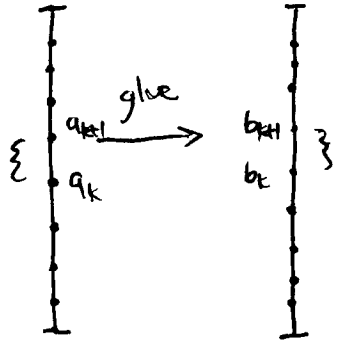

Fig. 13

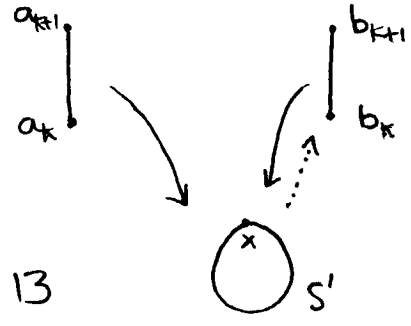

Foliations and the Topology of 3-manifolds

Outline of class 27

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

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

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

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

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

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

(5) D disconnects M, and $\pi_1(\text{eni})=1$ for exactly one i (say i=1).

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

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

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

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

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

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

If we take the arc transverse to each of the rectangles that we are gluing in pairs, it easy to see that each starts at L and then runs transverse to $\mathcal{F}_1$ tending to $R_\pm(\gamma_1)$. This path runs either with or against the normal orientation of $\mathcal{F}_1$, depending upon which side

of the picture the rectangle originates (i.e., lies in the cut) from. Because the space of leaves (when we throw away the $R_\pm$) is S^1, if we use (the point corresponding to) L as basepoint, we can partition the intervals, as before, into an infinite number of subintervals all wrapping around S^1 once as we march out to $R_\pm$ (Figure 14). If we do our gluings, as before, to respect this partitioning (i.e., after the identification with S^1, are gluings over each subarc is actually the identity), then we will always be gluing a leaf of $\mathcal{F}_1$ to itself (again, as in the corresponding case from last time). Therefore, $\mathcal{F}$ will have the same space of leaves as $\mathcal{F}_1$ (essentially, it has the same leaves, after throwing away $R_\pm$ (there could by components being glued to one another in $R_\pm$)). This proves the inductive step in this case.

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

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

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

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

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

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

Proof: Suppose not; suppose that for any n there is a compact leaf L_n of $\mathcal{F}$ and an i such that L_n meets $\mathcal{U}_i$ in more than n plaques. Using the principle that if you throw an

infinite number of things into a finite number of holes, you must have thrown an infinite number of things into one of the holes, we can assume that i=1 for all n. Now take a transverse arc α to the distinguished chart $\mathcal{U}_1$ (see Figure 16), and look at the $L_n \cap \alpha$. This consists, for each n, of more than n points. However, since each L_n is compact, this also consists of a finite number of points for each n - otherwise, these points of intersection would give an infinite discrete set in a compact leaf, which is impossible. But now we will inductively choose a sequence of compact leaves (which we will still call L_n) with the following property:

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

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

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

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

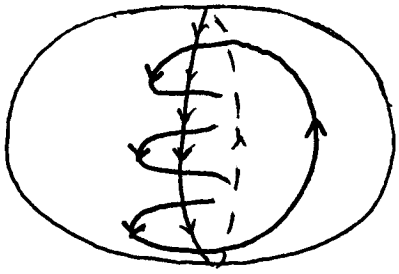

Fig. 1

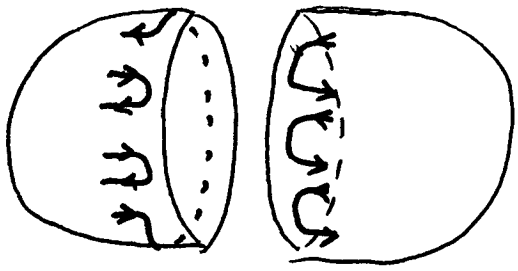

Fig. 2

Fig. 3

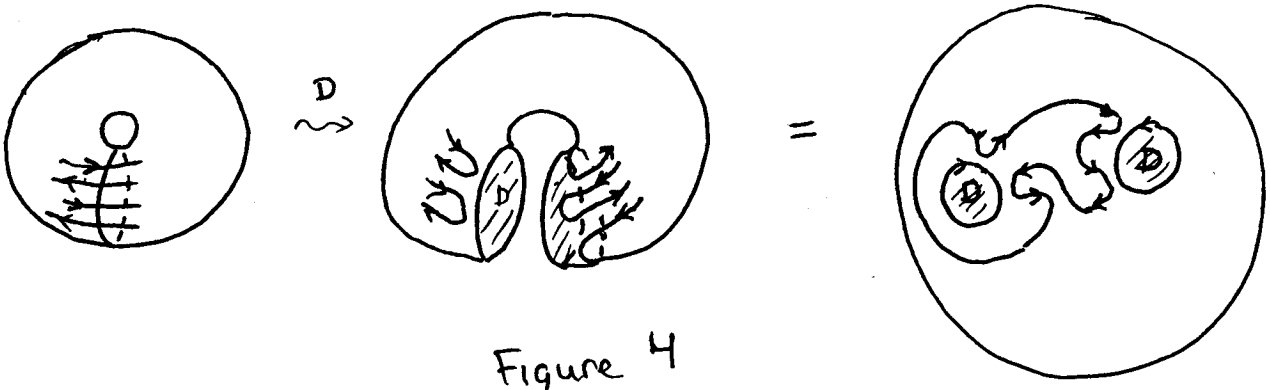

Figure 4

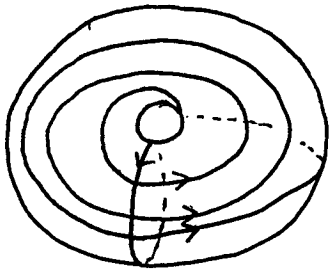

Fig. 5

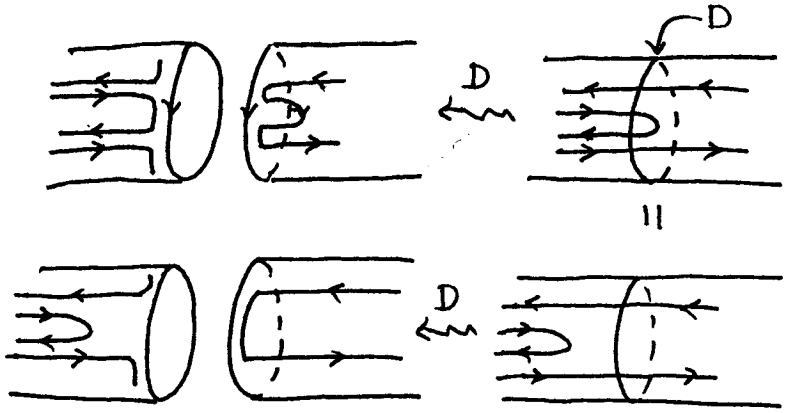

Fig. 6

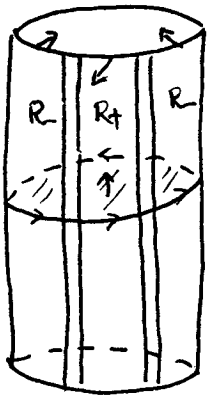

Fig. 7

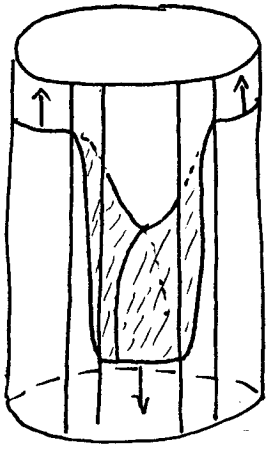

Fig. 8

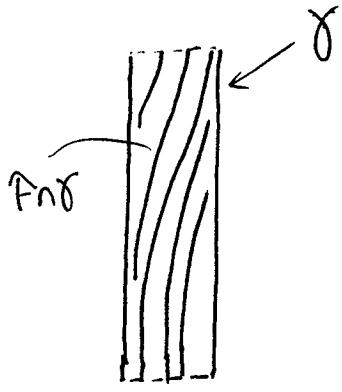

Fig. 9

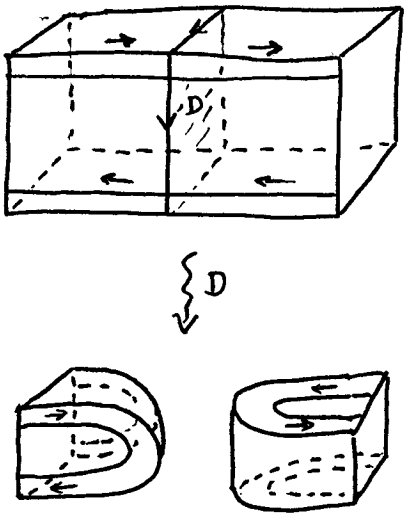

Fig. 10

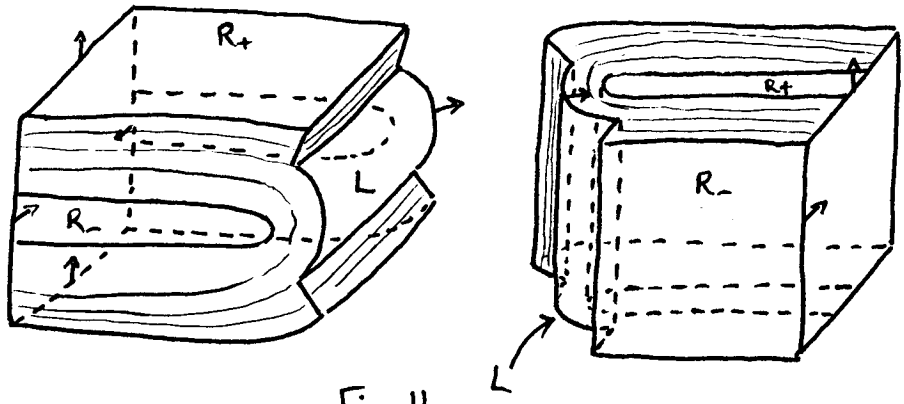

Fig. 11

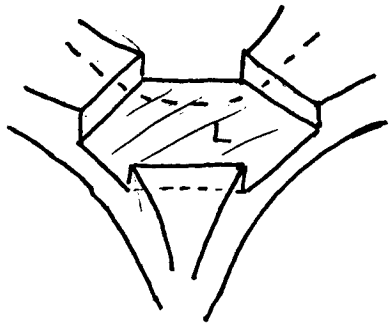

Fig. 12

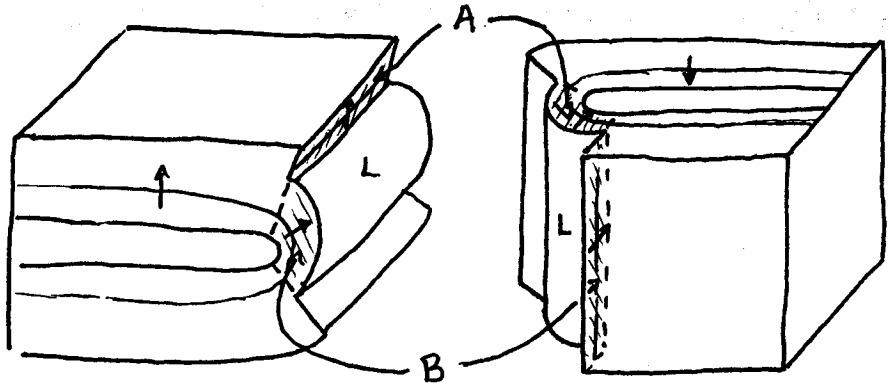

Fig. 13

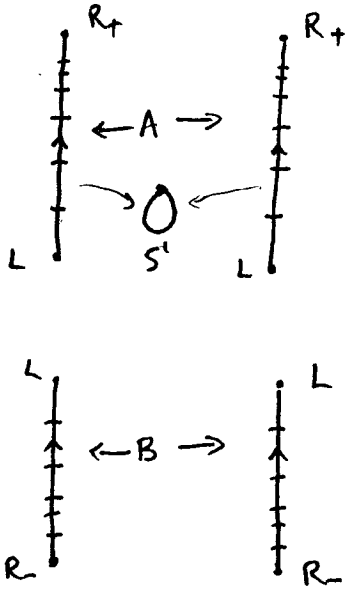

Fig. 14

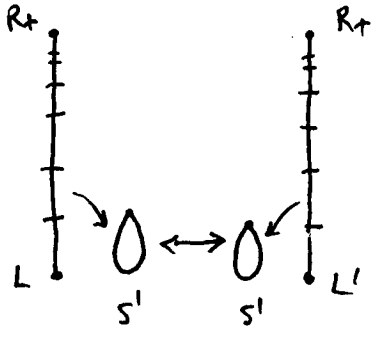

Fig. 15

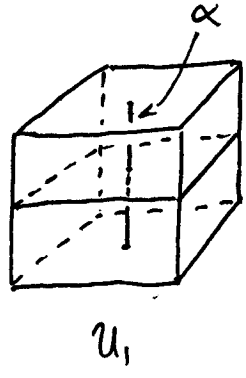

U_1

Fig. 16

Foliations and the Topology of 3-manifolds

Outline of class 28

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

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

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

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

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

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

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

Proof: The second part implies the first: if $\mathcal{F}$ has compact leaves L_n which are pairwise non-homeomorphic, then by choosing a point x_n of each, we get a sequence in the

compact M, so has a convergent subsequence (call it x_n) converging to some x (which by the above is contained in a compact leaf L). But the second part says that then for large n, $L \cong L_n$, so in particular are homeomorphic to one another, a contradiction.

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

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

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

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

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

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

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

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

Proof: Suppose not: then since $\mathcal{F}\cap\partial$M contains a compact loop (the boundary of the Seifert surface), and since we know that the set of loops is closed, we must be able to find a situation like Figure 5. We are assuming that the set of compact loops of $\mathcal{F}\cap\partial$M is not the entire boundary, so if we start walking transverse to the foliation (starting at a compact loop) we will find a last closed loop; beyond that we must have a non-compact leaf λ spiralling down on a compact loop γ. Call the leaf of $\mathcal{F}$ containing λ L, and the leaf of $\mathcal{F}$

containing γ S. Notice that we immediately have $\gamma \subseteq \overline{\lambda} \subseteq \overline{\ell}$, so $S \cap \overline{L} \neq \emptyset$, and consequently $S \subseteq \overline{\ell}$ (Exercise: just take any point of S and join it to a point of γ by an arc in S). Consequently, by (2), S is compact (or S=L; we will show why this is impossible after the rest of the proof).

Now we show that $\partial S = \gamma$; this follows from our work on Thurston norms! This is because the homology class $[\partial S]$ goes to 0 under the homomorphism $H_1(\partial M) \to H_1(M)$ (since it is clearly a boundary in M!). It therefore is represented by $n \cdot \ell$, where ℓ is the longitude of our knot (this is basically the definition of longitude). But by Alexander duality, S then represents n in $H_2(M)$; but since S is connected, and we showed that if a surface represents $\alpha + \beta$, it is the union of two disjoint surfaces each representing one of the classes, it follows that $n = \pm 1$. But since our hypothesis of no Reeb annuli in ∂M implies that all of the compact leaves in ∂M are parallel and coherently oriented (Figure 6), it then follows that ∂S is connected.

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

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

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

If $\overline{L} = L \cup$compact leaves, then L cannot limit transversely on itself. for if it did then if we look at a transverse arc α through such a limit point x, every point y of $B = \overline{L} \cap \alpha$ is a limit point of B. If $y \notin L$, this is clear (y <u>is</u> a limit point, and this limiting must be from the

transverse direction rather than from in the leaf containing it) while if $y \in L$, by joining it to x by an arc in L and looking at the normal fence over this arc, by the usual argument, we find arcs in L limiting down on this arc, giving, at the end y, a sequence of points in L limiting down on y. $\overline{L} \cap \alpha$ is therefore what is known as prefect. Of course, $\overline{L} \cap \alpha$ is also closed (in the interval α).

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

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

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

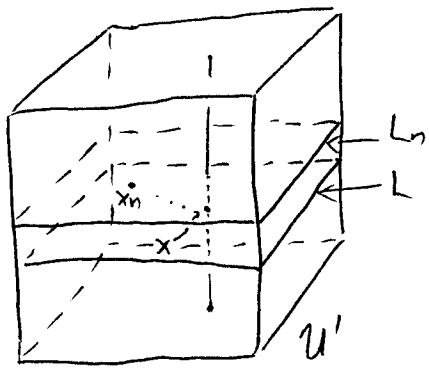

Fig. 1

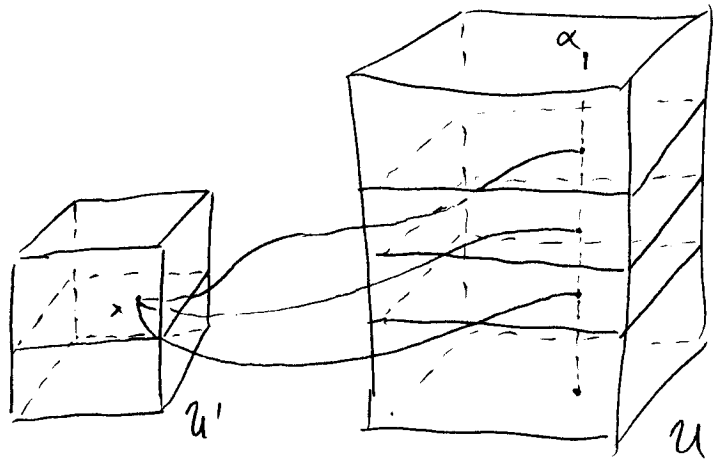

Fig. 2

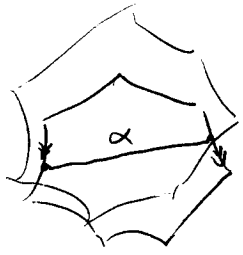

Fig. 3

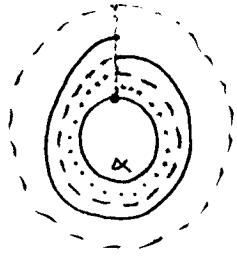

Fig. 4

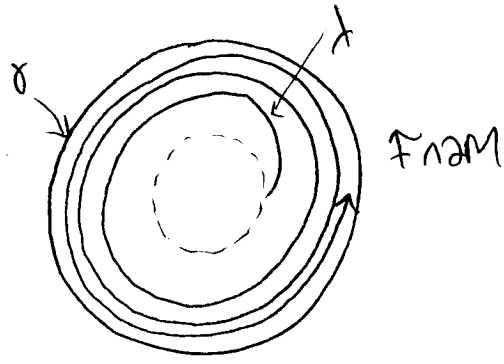

Fig. 5

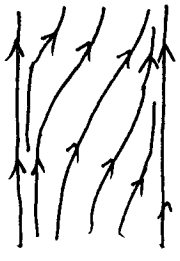

Fig. 6

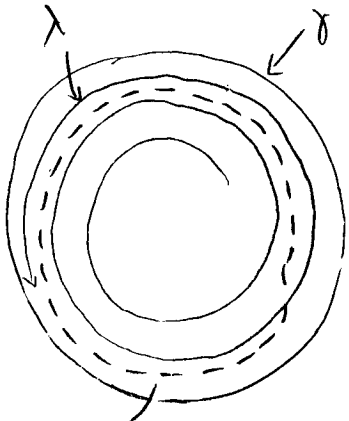

Fig. 7

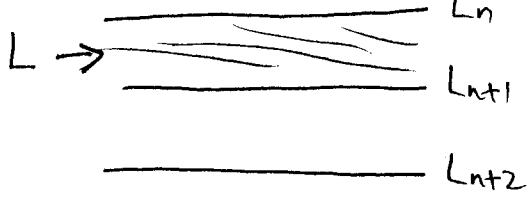

Fig. 8