Low Dimensional Topology Notes
June 30, 2006

Gabriel C. Drummond-Cole

June 30, 2006

1 Khovanov

[Good morning, a couple of announcements. The morning problem session will be on Khovanov's lectures. The afternoon lectures will be on Gordon's lectures. There is not going to be a long weekend, avoid the temptation to think there is.]

Last time I introduced $H^n = \bigoplus F(w(b)a)\{n\}$, where the shifts ensure that the idempotents are in degree zero. The ring is non-negatively graded.

Exercise 1 $(H^n)_0 = \prod_a \mathbb{Z}1_a = \bigoplus_a \mathbb{Z}1_a$.

Last time we assigned a bimodule $F(T)$ to a flat tangle T. This is an (m, n)-bimodule, meaning an H^m, H^n-bimodule.

Now it's time to add crossings. So for T we want to assign bimodules. A single crossing, we don't have to think, we simply resolve the diagram. We take the complexes $F(T_0) \to F(T_1)$. The differential comes from the cobordism which is creation of a saddle point. This induces a map of bimodules. We need to shift $F(T_1)$ by $\{-1\}$ to get this in degree zero.

In general, given an arbitrary tangle, you break it up into a product of tangles with at most one crossing and define $F(T) = F(T_k) \otimes F(T_{k+1}) \otimes \cdots$

Theorem 1 *$F(T)$ is an invariant of T in the homotopy category of graded bimodules (up to chain homotopy)*

This theorem says the complex does not depend on the representation so it's invariant under the Reidemeister moves.

The next step is to go to cobordisms of tangles and those will give morphisms of complexes of bimodules.

So T is a link with $n = m = 0$ then $H^0 = \mathbb{Z}$ and $F(T)$ is a complex of graded $\mathbb{Z}$-modules and then $H(F(T))$ is the homology from earlier, from lecture three.

A tangle cobordism connects two tangles thusly. It is an orient surface with corners S sitting in $\mathbb{R}^2 \times [0,1] \times [0,1]$. When you have an odd number of endpoints you somehow have to cheat, but in a link you always have an even number. Also, at some point you should add an orientation, and the corresponding shift $[x(T)]\{2x(T) - y(T)\}$ in the complex.

To pass to homomorphisms of bimodules, you want to shift $F(S)$ by 0 in the homological grading and $-\chi_S$ in the other grading. This is the same as the shift in a two-dimensional TQFT. You chop S into a bunch of pieces, and describe S by a sequence of tangle projections, planar diagrams. For instance you cauld have first a third Reidemeister move, and then saddle point moves and birth and death moves. This is a combinatorial representation of a tangle cobordism, and it's called a tangle movie.

Let S be a movie. Then to S we can assign $F(S)$, which is a composition of homomorphisms between two successinve tangles in the movie.

So for the Reidemeister three move you just have the canonical quasiisomorphism. For the next piece you have a change in topology. Take the corresponding map of bimodules $F(\tilde{T}) \to F(\tilde{\tilde{T}})$ and compose it with the identity on every side. Then just compose these in a number of ways. So it's a composition of quasiisomorphisms, saddles, births (unit) and deaths (trace).

What is the problem? We didn't assign a homomorphism to a cobordism, we assigned one to a decomposition of a cobordism.

Theorem 2 *If movies S_0, S_1 represent the same cobordism up to boundary fixing isotopy, then $F(S_0) = \pm F(S_1)$.*

Proof. First we have to find out which movies represent the same cobordisms. This was worked out by Roseman, Carter, Saito and is well-known. They made the list of movie moves that generate everything. There are quite a few of them, I'll write down only one. This is a Reidemeister II move, and then this is a Reidemeister III move. You could also do the move at the bottom and then apply the III move. Then you end up with the same tangle. That's one movie, this is another movie. The middle diagrams are different.

So you want to say the maps $F(S_0)$ and $F(S_1)$ are the same up to maybe an overall minus sign. The $F(S_0)$ is an isomorphism of bimodules. Then the degree of $F(S_0)$ is the same as the degree of $F(S_1)$, is zero. So $F(S_1)^{-1}F(S)$ is an isomorphism $F(T) \to F(T)$. This is because T is a braid. So then $F(T^{-1}) \otimes F(T)$ should be the ring H^n, or really the complex $0 \to H^n \to 0$.

We have $F(T) \to F(T)$ and then we can tensor the whole thing with $F(T^{-1})$ and get a map $F(T^{-1}) \otimes F(T) \to F(T^{-1}) \otimes F(T)$. This is $1 \otimes F : H^n \to H^n$. Call this map α. So α is invertible. It has degree zero. It is also a map of bimodules. What are the bimodule maps from $H^n \to H^n$. These are in bijection with central elements of H^n, left multiplication and right multiplication. So what about the degree zero part?

Exercise 2 *That's isomorphic to $\mathbb{Z}$ so that any such element is $m1$.*

So the map is ± 1. Hence $F(S_0) = \pm F(S_1)$.

You just play around with bimodules, and keep in mind that all of these are invertible. When you change topology it's only a little harder. The extra grading helps you. Everything is very rigid because you have control over both gradings. If you don't use the grading, the center of H^n is huge, it's isomorphic to H^* of the space of flags in $\mathbb{C}^n$ which fix the nilpotent operator with two size n Jordan blocks.

There are definitely -1 which are hard to get rid of. Scott Morrison has figured out how to get rid of these, he'll talk about it Monday.

This is again a two-functor, now projective, from the two-category of tangles and tangle cobordisms to the two-category of natural transformations between (n, m)-bimodules.

We can now restrict to links and link cobordisms to get the same theorem. Then $H(S)$ represents a particular homomorphism $H(L_0) \to H(L_1)$ up to a possible minus sign.

Bar-Natan just looks at categories $C(H_n)$, and he sees not just $(A, \mathbb{Z})$ where $X^2 = 0$, but the deformation $X^2 = t$, so a deformation to $(A_t, \mathbb{Z}[t])$, and everything works as before. His approach is more geometric. In particular he preserves the invariance under Reidemeister moves in a very geometric way. This allows one the fastest known way to compute homology on the computer.

You put L in thin position, meaning you want to minimize the number of crossings on every horizontal plane. Each time you, well, to the bottom tangle you assign $F(T)$ a complex of projective modules. If you see $0 \to P_a \to P_a \to 0$ you throw it out, simplifying as much as possible, and then going up a level corresponds to tensoring with another complex. And you go on. This is very efficient and now he can compute homology for very large knots. Things are labeled by crossingless matchings, which has $\frac{1}{n+1}\binom{2n}{n}$, so the main size of the link is the number of strands in a horizontal cross-section.

Jake Rasmussen found [unintelligible]for the deformation to $X^2 = t$. Now you can get $H_t(L)$ a module over $\mathbb{Z}[t]$. When you assign the theory to a diagram, to disjoint circles you assign the object $(A_t)^{\otimes k}$, tensoring over $\mathbb{Z}[t]$. Any finitely generated module over $\mathbb{Z}[t]$ decomposes into a torsion part and a free part $H_t(L) = Tor(L) + H'(L)$. Now E. S. Lee has shown that $H'(L) \cong \mathbb{Q}[X]$ in degree zero for a knot.

So $H(T)$ looks like a skyscraper, and around it is a bunch of torsion. If you quotient by torsion you get a smaller theory $H'(L)$. This is functorial under movie cobordisms of links. You have to shift $H'(L) = \mathbb{Q}[X]\{s(L) - 1\}$. This is the Rasmussen invariant.

So S is a function from knots to even integers and the theory has nice functorial properties. If you have a connected cobordism then the resulting map is nontrivial. This gives you a bound on the slice genus of the knot. The slice genus of L is the minimal gunus of a surface in the four-ball that bounds L. So you have $H'(\bigcirc) \to H'(L)$ which is nontrivial. The trivial knot starts in degree (-1) with 1 and then goes up by degree 2. So then the degree of the map is precisely $2g$ where g is the genus, but it must be at least the minimal genus of $H'(L)$, which is $s(L) - 1$. So the difference cannot be less than $2g$. So $g \geq |\frac{s(L)}{2}|$. The minimal cobordism satisfies this relation. In general it's hard to compute $s(L)$. You don't know where $s(L)$ starts. It's nontrivial. But when L is a positive knot, it all lies in the positive degree, and computing the position is easy. Then $s(K)$ turns out to be $n + 1 - c$ where c is the number of Seifert circles in the diagram and n is the number of crossings. So $g(L) \geq g_4(L) \geq \frac{n+1-c}{2} \geq g(L)$ so all of these are equalities.

This is the first algebraic proof of the Milnor conjecture that $g_4(T_{p,q}) = \frac{(p-1)(q-1)}{2}$. This was proved in 1991 by Kronheimer-Mrowka. This proof is due to J. Rasmussen. This is the first application we've discussed. There was a talk yesterday which talked about an application to Legendrian knots. To me I don't care if there are applications, if it has a nice structure. It's structurally very easy and also very rigid.

[The way you described the s-invariant is different.]

To me this is similar to, well, A was $H^*(S^2)$. Then A_t is the $SU(2)$-equivariant cohomology $H^*_{SU(2)}(S^2)$ where $\mathbb{Z}[t]$ is $H^*_{SU(2)}(\cdot) = H^*(\mathbb{HP}^\infty)$.

In the next lecture I would like to talk about generalizing the HOMFLY-PT polynomial to a trigraded cohomology theory and talk about the relation to Hochschild homology.

2 Gordon

The organizers have asked me to remind you that Monday is not a holiday. I also want to apologize for talking at the same time as Argentina Germany overtime. Background music will be provided at the tent, if you want to come drink beer or whatever. Maybe do math.

Remember that K is a hyperbolic knot, and we want to know when $K(\alpha)$ is not hyperbolic. So for the figure eight, $K(0)$ is a torus bundle over the circle and $K(4)$ contains a Klein bottle. In fact $K(1)$, $K(2)$, and $K(3)$ are Seifert fibered surfaces. Because the figure eight is amphichiral, you also get that the negative versions of these are also nonhyperbolic. So beside the trivial surgery which gives S^3 there are nine nonhyperbolic surgeries. This is the worst example known.

What we mean by attaching a two-handle along a knot K is, X is a three-manifold and K a knot in the boundary of X. Then there's a neighborhood of K which is just an annulus. Then two-surgery is just attaching $D^2 \times [-1, 1]$ with an attaching map $g : S^1 \times [-1, 1] \to N(K)$. So $X[K]$ will be the result $X \cup_g D^2 \times [-1, 1]$. So $\delta(X[K])$ is δX surgured along K, namely $\delta X - N(K) \cup D^2 \times \{\pm 1\}$.

Let F be a closed surface in S^3 with $S^3 = X \cup_F X'$ and let $K \subset F$ have the induced framing m. This is just the linking number of α on the edge of $N(K)$ with K. Push K off the surface in a normal direction and count the linking number. I claim if you do m-Dehn surgery is $X[K] \cup_\delta X'[K]$. Let me show you why this is true. What are you doing? So what is M-surgery? Look at the exterior of K. It's clear $M_K = X \cup_{F_0} X'$ where F_0 is $F - N(K)$. Now $K(m)$ will be $M_K \cup V$ where V is a solid torus and the boundary of the meridian disk is glued in along a curve of slope m.

So on the solid torus V on the boundary, you see the two curves α_1 and α_2. These are the curves that bound disks D_1, D_2 in V. Now α_1 and α_2 cut δV into two annuli A, A'. Then the annulus A is glued on here, and A' is glued on here. And D_1, D_2 cut V into two two-handles. I want to choose $A = \delta V \cap X$ and $A' = \delta V \cap X'$. Of course we want $A \subset \delta H$ and $A' \subset \delta H'$. Now $M_K \cup V$ is $(X \cup_{F_0} X') \cup_\delta (H \cup_{D_1 \cup D_2} H')$ which is $(X \cup_A H) \cup_\delta (X' \cup_{A'} H') = X[K] \cup_\delta X'[K]$. Now $\delta X[K]$ is $F_0 \cup D_1 \cup D_2 = \hat{F}_0$.

This is a general assertion, when you do the surgery corresponding to the framing of a knot on a surface you get a manifold with this nice decomposition. So we can try to use this to get an interesting class of examples.

Right. So we'll now take F to be a genus two Heegaard surface. Now X and X' are genus two handlebodies sitting in S^3. Let's get a definition that applies to curves on the boundary of a handlebody of genus two. Say $K \subset \delta X$ is primitive if and only if when you attach a two-handle along K you get a solid torus. This is equivalent to saying there's a properly embedded disk $D \subset X$ such that $|K \pitchfork \delta D| = 1$.

The reason it's called primitive, it's also equivalent to saying X represents a basis element in the free group $\pi_1(X)$.

So here's a primitive curve. That's not a very interesting one. The thing about handlebodies is that they're complicated objects. Luckily there are a lot of less obvious primitive curves, like this one. There's an exercise to prove that that's also a primitive curve.

[Can you show us what you did, you hid it when you drew it.]

I picked up the chalk like this. Maybe it's not primitive, that's not my problem, it's your problem. But there are a lot of primitive curves. There are infinitely many that bound disks. You get them all by handle slides or whatever you like. There is a disk here that intersects this guy in one point. I wouldn't recommend trying to find the disk. But you can prove in some other way that it's primitive. We're interested in our genus two surface in S^3. So I'll stick with this setup, and F is doubly primitive if its primitive on X and X'. I think this one is primitive in this handlebody I showed you but not in the other one. There are infinitely many, they've been completely classified. But suppose you have such a thing. Then $K(m)$, by this general discussion, is $X[K] \cup_\delta X'[K]$ is the union of two solid tori, so a lens space (also S^3 or $S^1 \times S^2$ but not here). So you just stay away from torus knots and satellite knots, and here you'll have hyperbolic knots with surgeries to lens spaces. This is the Berge construction. These doubly primitive knots have been completely classified as curves on the boundary of a handlebody, so thinking of them explicitly people call them the Berge knots. But the Berge conjecture, nobody knows any other knots with lens space surgeries. If K is a knot and $K(\alpha)$ a lens space for some α if and only if K is a Berge knot. There are some cases where some knot has more than one. For a given representation as a Berge knot there's a unique slope. Ozsvath and Szabo have some interesting results on this. This points up the facts that seem to emerge, that the knots that do have exceptional surgeries are somewhat simple.

Those are the Berge knots. Once you have this idea you can go to town with this, and make another definition. Say $K \subset boundary X$ is Seifert if when you add the two handle, $X[K]$ is a Seifert fibered space, either over the disk with two singular fibers or the Mobius band with one. Here's an example of a Seifert curve.

That, it turns out that if you attach a two handle along that curve you get the exterior of the trefoil knot, it's the Seifert fibered space over the disk with two exception fibers with index $(2, 3)$. The fundamental group, you can see a three manifold with $\pi_1 = \langle x, y, x^2, y^3 \rangle$, which si the fundamental group of the trefoil.

We can make another definition. Now we go back to the situation I'm erasing where K is sitting on F in S^3. We say it's primitive Seifert if it's primitive on X and Seifert on X'. There are lots of these although they haven't been completely classified. So what do you get when you perform the surgery corresponding to the framing? You get $X[K] \cup X'[K]$. This is a solid torus and a Seifert fibered space. So generically this will be a Seifert fibered space (occasionally a connect sum of lens spaces).

A bunch of examples of these were found by John Dean in his thesis. A hyperbolic knot in S^3 with Seifert fibered space surgeries don't have to be Dean knots, but most of them are.

[Question]

Tunnel number one knot exteriors can be gotten by adding a two-handle to a genus two handlebody.

If K is Seifert on both sides this gives you $K(m)$ a graph manifold or a Seifert fibered space (this essentially never happens). Generically it will be a graph manifold. And in particular, this common boundary, this torus, you see, it will be incompressible, in $X[K] \cup_T X'[K]$. Here's a way to find guys containing incompressible tori.

More generally you can try to ensure that the common boundary is incompressible in both halves. Even if you don't care what the pieces are, at least you have an incompressible torus.

A graph manifold is a union of Seifert fibered spaces glued together along tori.

In the generality I've given here, these may not be irreducible. You can get connected sums from cable knots and torus knots. But by the cabling conjecture, it will be true if you start with a hyperbolic knot.

Here's an example. Let me draw the knot first this time, the $(-2, 3, 7)$ pretzel knot.

We can put K on the surface of a two-genus handlebody. The first question is, what's the induced framing?

I'll leave this as an exercise, but I'll give you a hint,

Exercise 3 *Remember one way to think about m, there's a very natural way to find a surface here. This m, actually, it's the boundary slope of the surface with boundary. So how often does that boundary link the knot? It's a purely local consideration. Try to figure out how many times it links the orange. When the arrows are going the opposite direction, this does nothing. In the same direction, you have $7+3=10, 10\times 2$ is twenty.*

Then $K(20) = X[K] \cup_T X'[K]$ is Seifert-Seifert, a graph manifold.

That's another general philosophy. This is a construction.

Later we'll encounter this example again and be able to figure out all of thee nonhyperbolic surgeries. You have a hyperbolic knot K and a nonhyperbolic Dehn surgery. You have X hyperbolic and $K(\alpha)$ nonhyperbolic. It could be

1. S^3
2. $S^1 \times S^2$
3. non-prime
4. lens space
5. Seifert fibered space of type $S^2(q_1, q_2, q_3)$.
6. toroidal

I've given you methods to show you that these all occur. The S^3 never happen unless you take the trivial surgery (Gordon-Lueke). The $S^1 \times S^2$ is called property R and was proved not to occur by Gabai. The third case shouldn't happen, via the cabling conjecture.

The last three that do occur, do we have a full classification?

[Is it clear, this $(-2, 3, 7)$-pretzel knot example?]

The one side, it's obvious, you have a once-punctured Klein bottle and end up with an I-bundle over a Klein bottle. But it's harder to see on the other side.

This pretzel knot is a Berge knot in two different ways so has two lens space surgeries. It, along with the figure eight, has a lot of nonhyperbolic surgeries.

4. What about lens spaces? The Berge conjecture is not proven. Heegaard Floer homology says something; an earlier result says that if K is hyperbalic and $K(\alpha)$ is a lens space then α is an integer. One thing about this construction I described is that they always lead to integral Dehn surgeries. In some sense this is pointing in a very general way toward the Berge conjecture.

5. When do you get Seifert fibered spaces? The precise conjecture is, there are more ways, it's not clear what the conjecture is exactly, but $\alpha \in \mathbb{Z}$ should be necessary. That is pretty much wide open. You should also be able to say that the Seifert fibered spaces you get will be quite small, $S^2(q_1, q_2, q_3) >$ or $P^2(q_1, q_2)$. In some sense you have to solve this before you worry about the knots. It's not known whether you can get more than three.

6. There is one class of knots left, the toroidal case. What can you say about hyperbolic knots with toroidal surgeries? Again these are integral surgeries. In fact that's not true. Right off the bat, there is a very interesting class of examples, Eudave-Muñoz constructed $K(m/2)$ with K hyperbolic and $K(m/2)$ a toroidal (graph) manifold.

So already we have to look at nonintegral things here. This says how many times you have to intersect the meridian. Luckily, John Luecke and I proved that if K is hyperbolic and $K(m/\ell)$ is toroidal for some m/ℓ where $\ell > 2$, then $\ell = 2$ and the knot is in this Eudave-Muñoz class.

So we're just left with the integral surgeries. This is, what can you say about K_{hyp} and then $K(m)$ toroidal if m is an integer.

We've seen examples from genus two handlebodies. But it's going to be a much more complicated situation which hopefully will be involved some day.

Okay, I think that's a good place to stop.

3 Szabo

You may recall that we have struggled through the definintion for Heegaard Floer homology for three-manifolds. One thing you can do next is computations. You can do it for $S^3, S^2 \times S^1$, and lens spaces, but in general it's hard. You can study the problem of what happens when you have a knot $L \hookrightarrow Y$ in your three-manifold. You choose a longitude for the knot and perform surgeries on it, recording the information via a rational number. In fact, there's a notion of integer surgery. You can perform $+n$ surgery on K to get $Y_n(K)$. As it gets bigger you get more $Spin^c$ structure. We proved that $\widehat{HF}(Y)$ (you can write HF^+ here) then $\widehat{HF}(Y_n(K))$, it will not be the same, but there will be a long exact sequence:

$$\begin{tikzcd} & \widehat{HF}(Y_n(K)) \arrow[dr] & \\ \widehat{HF}(Y) \arrow[ur] & & \widehat{HF}(Y_{n+1}(K)) \arrow[ll] \end{tikzcd}$$

Now we are thinking of deleting the knot, and depending on how you fill in that manifold, you get maybe the three-manifold back or some of the surgeries.

So you could first look at almost a Heegaard diagram $(\Sigma_g, \alpha_1, \ldots, \alpha_g, \beta_1, \ldots, \beta_{g-1})$. I can build up the handlebody on one side and on the other side I can glue in disks and then eventually end up with a surface with torus boundary. So I can look at the α, β pair to correspond to Y; I can find an α, γ pair to correspond to Y_n and an α, δ pair corresponding to Y_{n+1}. All that changes is the last circle.

The holomorphic triangle construction gives this, it's similar to the handleslide argument.

[Grading question?]

The grading is not preserved under this map. The best you can do is with mod two grading. I could just make the surgery but these are integer surgeries. I could make Y_n by adding a two-handle and then there's a four-manifold, a cobordism, and then somewhere [unintelligible], so the bottom line is that the grading is screwed up.

At least you have some techniques that seem very ad hoc. If you're in a very nice closed family, where it's obvious they're related by long exact sequences then you can do this inductively. Let's look at $\Sigma(p, q, r)$. Then $\widehat{\Sigma(2,3,5)} = \mathbb{Z}$ in degree -2 and $\widehat{S^3} = \mathbb{Z}$ in degree zero. If you compute it for $\Sigma(2, 3, 7)$, you get something like $\mathbb{Z}$ and $\mathbb{Z} \oplus \mathbb{Z}$ in something like 0 and -1.

I can list all the homology three-spheres whose Floer homology is exactly the homology of S^3. For instance $\widehat{HF}(\Sigma(2,3,5)\# - \Sigma(2,3,5)) = \widehat{HF}(S^3)$. You can also take k each of these. These are the only examples we have of this.

Great, so what else? Even in these long exact sequences we see cobordisms. If I have a compact oriented smooth four-manifold, what kind of invariant can we associate to that.

All of the old four dimensional invariants had versions for three manifolds. We were trying to do the reverse, starting from the three-manifold invariant. To construct that you have to think about $Spin^c(W)$. This uses the general definition. These can be identified with $H^2(W)$. So fixing a $Spin^c$ structure you can find a map to identify $\hat{F}_{W,s} : \widehat{HF}(F_1, s|_{Y_1}) \to \widehat{HF}(F_2, s|_{Y_2})$.

So this involves a Morse theory picture and some Kirby calculus. It's complicated. It looks like it depends on the decomposition but then it turns out it only depends on the smooth structure. In the examples where we can compute it it seems like the Seiberg Witten invariant.

Okay, so we still have time. Instead of going up in dimension we'll go down. Look at $K \hookrightarrow Y$. This is due to Jake Rasmussen and independently $Ozsvath - Szabo$. In this setting, instead of having one basepoint, I can choose two, in two different components. I can still try to do the same definition, almost the same. So

$$\delta x = \sum_{y \in T_\alpha \cap T_\beta} \sum_{\substack{\phi \in \pi_2(x, y) \\ Mas(\phi) = 1}} \#\mathscr{M}(\phi)/\mathbb{R} \cdot y, n_z(\phi) = 0 = n_w(\phi).$$

So the two basepoints give us the knot itself. One basepoint give a trajectory through the minimum and maximum. When you have two such points, then looking at both trajectories you can recover the knot. How to get from a two-pointed Heegaard diagram to a knot.

A slightly different definiton is, I will find my knot. Somewhere are my α and I can cut them out and the surface is still connected, and I can connect the two points by an arc. Then I can push it into the handlebody corresponding to the alpha. I can also push into the opposite handlebody in the complement of the βs This is the same thing, if you think about it.

First you have to show that you can find a Heegaard diagram for any knot. You have to think about stabilizations and so on.

Which of these basepoints do you use for the $Spin^c$ structure?

It doesn't matter if the knot is homologically trivial. There is a relationship between two different $Spin^c$-structures if it is not homologically trivial.

I'll avoid this by looking in S^3 where there is only one $Spin^c$ structure.

For the purposo of the lecture then let's focus on $K \subset S^3$. This is not the only knot I can draw but it is an easy knot for me to draw.

We would like to construct a corresponding Heegaard diagram. I can prove this, but it might be better to start from a given projection. I will choose one edge of the knot to be special. Then I draw the picture forgetting about over and undercrossings. This is an immersion in the plane, but it lies in $\mathbb{R}^3$ as well. Thicken it there to get a handlebody. So now we have a Heegaard decomposition of S^3. I have to choose the α and β circles, which I do in the following way. Around the special point in an edge, I choose the circle parallel to the knot as a β. I draw this circle around a thickened crossing that looks like this, matching the over and under. I just constructed all of the β. Here I have two regions around the special edge. One I forget. Every other region gives us a contour. You have to see that these are linearly independent and they bound disks on the outside. That is clear. You should think that the surface has depth. This is a Heegaard diagram for S^3 but not yet a two-pointed one. So now I choose z here and w here. Well, I have to connect these two points in the complement in the α circles. That is easy. To do the other point, I have to go away from the β circles. If you try to follow the original knot, you are forced to follow along the original knot. You will get exactly the original knot pushing this inside the β handlebody. When I get to here I have to climb to the back. When I look in the knot inside S^3, then I really get this picture.

So then for any knot I am able to find a Heegaard diagram. I am trying to do this as combinatorially as possible. So okay. The generators. Let's write down the generators in $T_\alpha \cap T_\beta$. So let's look at this picture. This was the β circle, but there are also the α circles. So we get four intersection points here, so shorthand, for these points, we can record a point in each of the four corners of a region around a crossing. You are looking at intersection points in the symmetric product. I can start with the special circle which only intersects one other circle. So I have to choose that point. This seems very combinatorial. The generators are in one to one correspondence with what are called the Kauffman states.

The idea is for every double point, choose one of these corners. It seems like I have 4^n choices, but there is another condition, which is that every α circle must have one point on them, exactly. These points correspond to regions. So every region must have exactly one corner, except the region on the two sides of the marked point have exactly one corner. So I have n double points, $n+2$ regions, and now I leave out the other one adjacent to the marked point. I can't use any corner adjacent to this bad region. So Kauffman states don't originally have to do with Heegaard diagrams. But mark each other region one time. These generators are in one to one correspondence with Kauffman states.

In that example you are welcome to find all the Kauffman states and you will find three of them.

This will eventually compute something for us. The generators are exactly the Kauffman states. To better understand the chain complex, I'm going to write down some definitions. This will be local and correspond to corners. So around a double point, fix an orientation. I want a number. These will correspond to $\pm 1/2$ and then 0 for the corners as indicated:

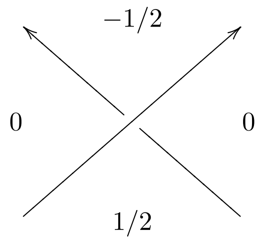

or

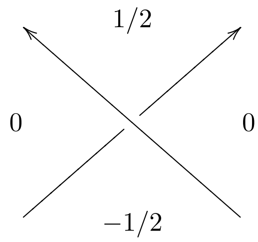

for A and then

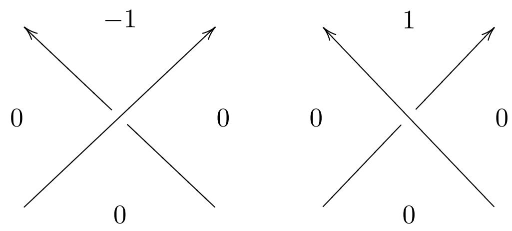

for M.

Then $A(x)$ is $\sum A(x_i)$ where $\underline{x}$ is a Kauffman state $(x_1, \ldots, x_n)$. Then $M(x)$ is similarly $\sum M(x_i)$. So why do we do this.

What we did now, I can write the Kauffman state somewhere in the board, so that if $x \in T_\alpha \cap T_\beta$, then there is a corresponding Kauffman state and $i(x)$ will be A of that state and $j(x)$ will be M of the state. So I can write $C_{i,j}(K)$ which is generated by all the intersection points with appropriate i and j.

Lemma 1 *$x, y \in T_\alpha \cap T_\beta$ with $\phi \in \pi_2(x, y)$ then $n_z(\phi) - n_w(\phi) = i(x) - i(y)$.*

This is just a calculation.

Lemma 2 $Mas(\phi) = 2n_w(\phi) = j(x) - i(y)$.

If you have a holomorphic disk like that, the formula of Lipschitz and the other lemma tell you this. So the bounding map, we get out of this, $C_{i,j}(K) \to C_{i,j-1}(K)$. But it's harder to compute the boundary maps themselves. One thing you can do here is fix i and just look at the chain complex above it, and you can compute the Euler characteristic. So you can get $H_{i,j}(K)$. You can take $a_i = \sum(-1)^j \mathrm{rk}\, H_{i,j}(K)$. This is really easy because you are just doing computations on the Kauffman states. This was to come up with a combinatorial description of the Alexander polynomial. In fact, what you see is that this is nothing other than the Alexander polynomial, $\sum a_i T^i = \Delta_K(T)$.

This homology is the double-graded homology where i is the homological grading and j is the Maslov grading. So the Euler characteristic is uninteresting, it's just the Alexander polynomial.

Alternating knots are really easy. You can use these two lemmas to show that for an alternating projection, $C_{i,j}(X)$ will all lie on a line, which intersects the $i = 0$ axis on $sgn(K)/2$. Then all of the boundary maps are zero. So this homology does nothing other than compute the Alexander polynomial and the signature of the knot. This is good and bad news. It's a little bit of both. It seems like a random construction. When you would like to compute the Heegaard Floer homology for surgeries on the knot, when you write down the chain complex when n is large, you end up with similar chain complex for $\widehat{HF}(S^3_n(K))$. In another sense it's uninteresting because for alternating knots it gives something not really interesting.

So the question is, when is this interesting. I should have said that the Floer homology is easy for the unknot, a single $\mathbb{Z}$ in degree $0, 0$. If we have a nontrivial knot with $\Delta_K = 1$, can we compute the Alexander polynomial? So here is one. Let's hop I can do this.

What I want to draw is the Kinoshita-Terasaka knot. This is one of those knots with trivial Alexander polynomial. This is a test example. If the Floer homology is the same as for the unknot, it is maybe an uninteresting theory. If it's the same then maybe that's more interesting. You have a chain complex, you have to understand some boundary maps, and you get some homology zero which is zero outside $[-2, 2]$. You get two $\mathbb{Z}$ in each of these, and then something in the center which is bigger than $\mathbb{Z}$. It's hard to compute, easier at the boundary. For a general knot you have the notion of the genus, and it's relatively easy to show that outside $[-g(K), g(K)]$ you always get zero homology. In all the cases we looked at it's always nontrivial on the levels $\pm g(K)$. This is now

Theorem 3 *The homology* $\bigoplus H_{g(k),j}$ *is nonzero*

So this detects knottedness. You can use this to prove

Corollary 1 *If* K *is a knot in* S^3, *then if* $S^3_n(K) = S^3_n(U)$, *where* U *is the unknot, then* K

is in fact the unknot. In that form for $n = 0$ this is due to Gabai. For ± 1 this is due to Luecke, Gordon, and the general case is due to Kronheimer-Ozsvath-Szabo.

There are other applications but my time is completely done.

Maybe I could say something posthumously. It's interesting to relate this to Khovanov homology. You can just compute it for any small knot. We don't have a combinatorial definition. We don't know how to compute the boundary map. So how can you find a more combinatorial definition which does the same thing.

4 Hutchins

This is okay? All right, so I'm going to talk about embedded contact homology. This is based on joint work with Michael Sullivan (in Geometry and Topology, 2006), so you can read it on your laptops while I'm talking. Then some of it is joint with C. Taubes. The idea of the project is that Taubes proved the Seiberg Witten = Gromov theorem, that the Seiberg Witten invariant of a closed symplectic 4-manifold X equals a certain count of embedded (with some multiply covered torus components) of embedded J-holomorphic curves in X.

The goal is to describe the Seiberg-Witten Floer homology of a contact 3-manifold Y in terms of J-holomorphic curves in $\mathbb{R} \times Y$ (which has a natural symplectic structure). There's also a version for a mapping torus. The homology will be $C_*(Y, \lambda, \Gamma, J)$.

The setup, Y^3 is a closed oriented three-manifold, λ is a contact form on Y, a 1-form with $\lambda \wedge d\lambda$ always positive. Then $\xi = \ker \lambda$ and then the Reeb vector field R which is in the kernel of $d\lambda$ and is normalized so that $\lambda R = 1$. We pick a homology class $\Gamma \in H_1(Y)$ and choose an almost complex structure J on $\mathbb{R} \times Y$. There are some standard assumptions,

- J is $\mathbb{R}$-invariant
- $J(\delta_S) = R$
- K sends ξ to itself, with a positive rotation

Then the embedded contact homology is the homology of a chain complex C_* over $\mathbb{Z}$ and a generator of the chain complex is a finite set α of pairs $\{(\alpha_i, m_i)\}$ where

- the α_i are distinct embedded Reeb orbits (some definitions of Reeb would allow you to go around multiple times, this is only once) and m_i is a positive integer (the multiplicity).
- Then the total homology class is equal to γ so $\sum m_i[\alpha_i] = \Gamma \in H_1(Y)$.
- There's one more technical condition, if α_i is hyperbolic then m_i is 1. If the rotation has real eigenvalues it's called hyperbolic.

I haven't told you the grading, let's put that aside. The differential counts J-holomorphic curves $C \subset \mathbb{R} \times Y$ such that

- C has positive ends at covers of the the Reeb orbits α_i with total multiplicity m_i, so that as the $\mathbb{R}$-orbit goes to ∞, you are going to some multiple, like you could do a double cover, and have another coming in like this, and the total multiplicity is m_i.
- C has negative ends at covers of β_j with total multiplicity m_j. Here this is the differential from $\alpha = (\alpha_i, m_i)$ to $\beta = (\beta_j, n_j)$.
- The most important, and technical part, is that $I(C) = 1$. This I is the embedded contact homology index. The important part is that if you have $I(C) = 1$ then C is embedded except for multiple covers of $\mathbb{R}$-invariant cylinders. You might have a multiple cover of a Reeb orbit cross $\mathbb{R}$. Further, the embedded piece has the usual index of symplectic field theory, Fredholm index 1. So it makes sense to count these things.

Those are the things we want to count. I haven't said anything about the topological constraints of C. In the examples I know the curves are all genus zero, but they could be anything, I suppose.

The claim, that I'm writing the paper with Taubes about, is $\delta^2 = 0$. Hopefully the methods will generalize to show that the homology ECH_* will depend only on Y, ξ, Γ.

What could we hope to do with this? Are there any questions?

[Does it really depend on ξ?]

I expect that it doesn't really, in a way that I'll state

Conjecture 1 *This is the big conjecture. $ECH_*(Y, \lambda, \gamma)$ is isomorphic to the Seiberg Witten Floer homology $\check{HM}(-Y)$ which is also conjecturally isomorphic to $HF^+(-Y, \iota_\lambda(\Gamma))$.*

Here $\iota_\lambda : H_1(Y) \to Spin^c(\pm Y)$.

Why would we want more versions of Floer homology, isn't one enough? The Weinstein conjecture says that for all Y^3 closed oriented with contact form λ there exists a closed Reeb orbit. The big conjecture implies the Weinstein conjecture, because some results of Kronheimer and Mrowka indicate that if $C_1(s)$ is torsion (s a $Spin^c$ structure) then $\check{HM}(Y, s)$ is infinitely generated.

If there are no Reeb orbits then $ECH_*(Y, \lambda, \Gamma)$ is $\mathbb{Z}$ for $\Gamma = 0$, and is zero otherwise, which would contradict this.

[Can you say there are infinitely many?]

There are various theorems where under certain conditions, you could have either two or infinitely many.

There's some additional structure on this. For example, in here we're sort of imitating the other Floer homologies

- Ther is a map $U : ECH_* \to ECH_{*-2}$ which is defined by counting $I = 2$ curves that pass through a distinguished point z in $\mathbb{R} \times Y$. You set the index equal to two so you get a two dimensional moduli space of these things.
- There's also a canonical element which we've already seen, called $c(\lambda)$, which is the homology class of the empty set $[\emptyset]$. This is a cycle because a holomorphic curve has a positive end. The Floer theories have an invariant of contact structures which corresponds to $-Y$. This is as canonical as this could be.

Let's do an easy example. It's been a long week.

Take $Y = S^3$. There's a standard contact structure so that the Reeb orbits are the circles of the Hopf fibration
$$\begin{tikzcd} S^1 \arrow[r] & Y \arrow[d] \\ & S^2 \end{tikzcd}$$
I think the form is $\lambda = \frac{1}{2}(x_1 dy_1 - y_1 dx_1 + x_2 dy_2 - y_2 dx_2)$.
Then the Reeb orbits are the fibers or covers thereof.

We can perturb to get two Reeb orbits a and b, both elliptic. They could be long, but if you say things correctly those don't matter. Here of course Γ is zero. A generator is $a^m b^n$. Then m, n are nonnegative integers.

When you compute the index, which I didn't define, you get $I(a^m b^n) = (m + n)(m + n + 1) + 2m$. What is that, let's draw it.

	12			
	6	14		
$n = 1$	2	8	16	
$n = 0$	0	4	10	18
	$m = 0$	$m = 1$		

So then you have

$$\begin{tikzcd} \mathbb{Z} \arrow[rr, bend left, "\overset{U}{\cong}"] & 0 & \mathbb{Z} \arrow[rr, bend left, "\overset{U}{\cong}"] & 0 & \mathbb{Z} \arrow[rr, bend left, "\overset{U}{\cong}"] & 0 & \mathbb{Z} \end{tikzcd}$$

Okay, for a harder example let $Y = T^3_{\theta,x,y}$ Then $\lambda_n = \cos(n\theta)dx + \sin(n\theta)dy$ and $R_n = \cos(n\theta)\partial_x + \sin(n\theta)\partial_y$. This is the joint work with Michael Sullivan. We calculated $ECH_*(T^3, \lambda_n, \Gamma)$ is $\mathbb{Z}^3$ if Γ is zero and $* \geq 0$ and otherwise 0. This is the untwisted version. For $n = 1$ the empty set is a homology generator and otherwise it's zero. In the twisted version you get a nonzero element for every n. If that disagrees with the other Floer theories, please let me know.

Okay. Any more questions? Well now, I need to get a little bit technical. So you can dim the lights in the back of the room. I want to prove that $\delta^2 = 0$. So you have to do some gluing, of u_+ to u_-. The lower generator of the index one curve u_+ is the same as the upper generator of u_-. The only nontrivial case is when γ, the Reeb orbit, is elliptic. So u_+ has negative ends at covers of γ with multiplicities $a_1, \ldots, a_k$. Then u_- has pasitive ends at covers of γ with multiplicities $b_1, \ldots, b_\ell$. I know that $\sum a_i = \sum b_j$. The numbers don't have to match up, and when the sum is bigger than one, none of the them match up. The multiplicities don't match, so how are you going to glue them together?

You cross γ with a finite length cylinder $[-R, R]$ and glue in a branched cover of the cylinder. In general it can be disconnected, but I'll draw the connected case. This is like plumbing, so you have to insert some kind of piece, and I claim that one can glue like this. So how does this work?

Here $\mathscr{M}_R$ is the moduli space of branched covers, and over this is the obstruction bundle $\mathscr{O}$ with rank $2k + 2\ell - 4$. Then there's a section ψ which measures the obstruction to gluing.

$$\begin{tikzcd} \mathbb{R}^{2k+2\ell-4} \arrow[r] & \mathscr{O} \arrow[d, shift right] \\ & \mathscr{M}_R \arrow[u, shift right, bend right, "\psi"'] \end{tikzcd}$$

Here if ψ is zero you can glue. So $C \in \mathscr{M}_R$. then $\mathscr{O}_C = coker(D_C)$, and when $\psi(C) \cong \prod_{coker} \bar{\delta}(u_+ \#_R C \#_R u_-)$. Then the number of gluings is $\#\psi^{-1}(0)$. But to understand the zeros you need to understand what the section is doing at the boundary.

Theorem 4 *There is a combinatorial formula for the number of gluings. It depends only on* a_i, b_j*, and* θ *which, since this is an elliptic Reeb orbit, it has some return map which is a rotation* $2\pi\theta$*. The structure of the combinatorial formula is that it's a sum over at least one, some labelled trees, of certain positive integer contributions.*

I don't have time to state the full formula so I'll just state an example so you can see the kind of crazy thing you have to deal with.

[Why is m one for hyperbolics?]

The number of ways to glue is zero in the larger cases because they cancel with one another. Here we will have to glue along multiple elliptic things, to show $\delta^2 = 0$.

So let $k = 2, \ell = 3$. Then there are seven trees which may or may not contribute.

$$\frac{\lceil a_1\theta \rceil}{a_1} \le \frac{\lceil a_2\theta \rceil}{a_2}, \frac{\lfloor b_1\theta \rfloor}{b_1} \ge \frac{\lfloor b_2\theta \rfloor}{b_2} \ge \frac{\lfloor b_3\theta \rfloor}{b_3}$$

In these trees you have to assign weights, which give weights on the internal edges, and they only contribute if the edges have positive weights.

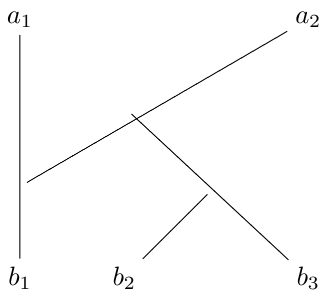

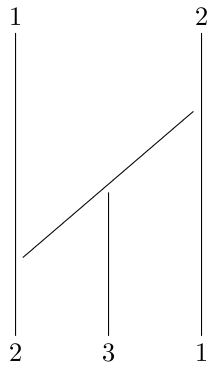

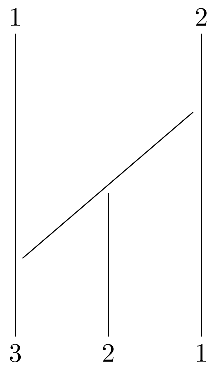

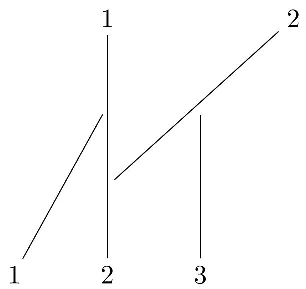

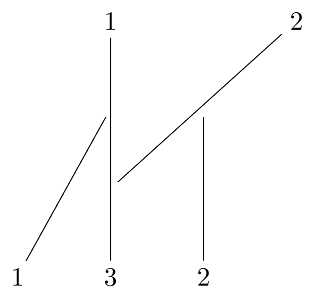

and two others.

In ECH situations $a_1, \ldots, a_k$ and $b_1, \ldots, b_\ell$ are determined by the condition $I(u_+) = I(u_-) = 1$ and then only one tree contributes with contribution one. The boundary of the moduli space has one point for each configuration, and that's why $\delta^2 = 0$. If you believe this is an analogue of Gromov is Seiberg Witten then this makes sense.

[For what three manifolds can we compute this?]

It's not too hard for a Seifert fibered space, I think. I also think for the T^3 case those methods could be generalized for torus bundles.

In general I don't know if a combinatorial formula would work. There might be something with open books.

[Something about overtwisted contact forms]

Whichever one has the simplest Reeb orbits would be the best.

[Does this theory give insight into what it means for $\emptyset$ to be a boundary?]

The usual argument that the contact homology vanishes for the overtwisted structure is roughly something like the Reeb orbit bounding the overtwisted disk. Some argument would show that for an overtwisted structure the canonical class is zero. If it's not overtwisted, I don't really know.

This is a fantastic theory where the empty set provides the most important invariant.