Low Dimensional Topology Notes
July 13, 2006

Gabriel C. Drummond-Cole

July 13, 2006

1 Etnyre

Okay, so, remember, yesterday we did a couple of things. If you start with a taut foliation and perturb it into a contact structure, so you can construct a weak semi-filling, so you get a tight contact structure. So we want now to construct the symplectic caps.

I'm going to start with the main theorem about the caps.

Theorem 1 *Eliashberg, Etnyre*
If (X, ω) is a compact symplectic manifold which weakly fills (M, ξ), then there exists a closed symplectic manifold (X', ω') into which (X, ω) embeds (as a symplectic manifold).

So we want to see how to construct the caps. I'm going to go over the proof I know better, which is the one that I figured out. Eliashberg's proof is different and gives you slightly different information.

I'm going to pretend that there's only one boundary component, but you'll see that the method would work for each boundary component seperately.

Definition 1 *Let Σ be an oriented compact surface with boundary. Let $\phi : \Sigma \to \Sigma$ be a diffeomorphism equal to the identity near the boundary.*

Then we can look at the mapping cylinder of ϕ, $T_\phi = \Sigma \times [0, 1]/(x, 0) \sim (\phi(x), 1)$.

For each boundary component, you get a torus. Let $M_{\Sigma,\phi} = T_\phi \amalg_{|\delta\Sigma|} (S^1 \times D^2)/\sim$, where we glue the solid tori to δT_ϕ so that $\{pt\} \times \delta D^2$ maps to the interval direction and $S^1 \times \{pt\}$ goes to the boundary of Σ.

So all of the pages in this picture come to the same core, that's why it's called an open book.

Exercise 1 *If L is the cores of the solid tori, then $M_{\Sigma,\phi} \backslash L$ fibers over S^1.*

Definition 2 *The fibers are called pages, and L is called the binding. An open book decomposition of M is a diffeomorphism to some $M_{\Sigma,\phi}$ as above.*

This is a weaker definition than we would usually use.

[This should be called a Rolodex.]

The terminology of open books is fairly well established. I don't know whether they had Rolodexes then. [Ed.: the Rolodex was first marketed in 1958, the term "open book" was introduced in 1973]

All three-manifolds admit open book decompostions.

Exercise 2 *Prove this. Hint: think about branched covers and braids.*

An open book Σ, ϕ for M is said to be compatible with or support a contact structure ξ if there exists a contact form α for ξ such that

1. $\alpha(TL) > 0$.
2. $d\alpha|_{page}$ is (up to sign) a volume form on Σ.

Theorem 2 *Thurston-Winkelnkemper*
Any open book supports a contact structure.

Given (Σ, ϕ) supporting (M, ξ), a positive stabilization of the open book is the open book with

1. page $\Sigma' = \Sigma$ with a one-handle.
2. $\phi' = \phi \circ D_\gamma$ where D_γ is a Dehn twist.

Exercise 3 $M_{(\Sigma,\phi)} \cong M_{(\Sigma',\phi')}$. *and the open books support the same contact structure.*

Theorem 3 *Giroux*
There is a bijection between oriented contact structures up to isotopy and open book decompositions up to positive stabilization.

This has been key to a lot of progress in contact geometry and topology in the last five years.

There's an analogous statement for higher dimensions, but it's more technical.

Next I want to see how surgery interacts with open books. So given (M, ξ) and an open book (Σ, ϕ) supporting it, let γ be a simple closed curve on the page of the open book.

Let $\mathscr{F}$ be the framing induced by the page, and let M' be obtained from M by $\mathscr{F} \pm 1$ Dehn surgery on γ.

Exercise 4 *Show an open book for M' is $(\Sigma, \phi \circ D_\gamma^\pm)$. It's a little strange that if you do a framing plus one surgery you get the negative Dehn twist.*

I'll givee you a hint in the form of the picture. This also helps show that every three-manifold is obtained by surgery on a link. I'm going to cut open along the page that contains γ. If I remove kind of a square neighborhood of γ, if I glue back everything except the square neighborhood, you get everything in M with the torus deleted. That's the same as in M'. If you stare at it long eonugh you'll be able to see it.

Now we're going to pay attention to the contact structure, so we'll want γ to be Legendrian. So we have this fact. If γ is nonseperating on the page, then we can isotope the page a little to make γ Legendrian and the contact framing agrees with the page framing.

This is a very nice result that allows us to do what we just did in the topological setting in the contact setting. So let (M', ξ') be obtained from (M, ξ) by Legendrian surgery. This is Dehn -1 surgery on the curve.

[What is the point of nonseperating?]

To make it true?

Okay, so the fact is that (M', T') is supported by $(\Sigma, \phi \circ D_\gamma)$. So given a symplectic (X, ω) filling (M, ξ) with open book (Σ, ϕ) and γ on a page, we attach a symplectic 2-handle to (X, ω) to get (X', ω') with $boundary(X', \omega') = (M', T')$ then we get that an open book for (M', T') is $(\Sigma, \phi \circ D_\gamma)$.

We want to build our caps, but to do this we need some facts about the mapping class group of a surface. If Σ is a surface with one boundary component then any diffeomorphism ϕ of Σ up to isotopy can be written as the composition of a bunch of twists $D_c^m \circ D_{\gamma_1}^{-} 1 \circ \cdots \circ D_{\gamma_n}^{-1}$ where c is a curve parallel to the boundary and γ_i are simple closed curves.

Okay, now we start to construct our caps. So given (X, ω) weakly filling (M, ξ), and (Σ, ϕ) an open book supporting (M, ξ), and assuming Σ has one boundary component, well, how can we do that?

[Stabilization?]

Exactly. If you do positive stabilization, you can connect boundary components and make sure there's only one boundary component. Also assume that ϕ is written as above, as $D_c^m \circ D_{\gamma_1}^{-} 1 \circ \cdots \circ D_{\gamma_n}^{-1}$.

Now, using the idea above, we can Legendrian "realize" all the γ_i, then construct (X', ω') by attaching 2-handles to (X, ω) along γ_i. So $\delta(X', \omega') = (M', \xi')$ with open book (Σ, ϕ'), where, well, ϕ was the composition of all of these Dehn twists. A Dehn twist along γ_n gets rid of $D_{\gamma_n}^{-1}$ so we get $\phi' = D_c^m$.

That's good enough for some people but I still don't see what's going on. Attach more 2-handles so that $\phi' = D_c^m \circ D_{\delta_1} \circ \cdots \circ D_{\delta_{2g}}$, where g is the genus of Σ and the δ_i form a chain of curves that serve as a basis, so that $\delta_i \cdot \delta_{i+1} = 1$.

Exercise 5 *if the genus of Σ is g, then M' is $\frac{1}{n}$ surgery on g-trefoils.*

Here's a fact. If (X, ω) is a weak filling of (M, ξ) and M is a homology sphere then we can slighly perturb ω so that (X, ω) is a strong filling. Now that I remind you of this fact, can anyone notice anything about M'?

[The M' is a homology sphere.]

Yes. Great, so, (X', ω') is a strong filling of (M', ξ') with (X, ω) embedded in it. So now I only have to tell you how to construct caps for strong fillings. But it's not that hard just to construct a cap for a strong filling. So let's just do that.

So now we can stabilize the open book for (M', ξ') so that the page looks like, well, the monodromy has a bunch of Dehn twists around the boundary curve, and then a chain around the genus. When we stabilize, we add links to the chain when we add more genus.

I can arbitrarily pick any open book for M, so I'm not affecting X here. So if we stabilize (Σ, ϕ') enough, then after adding positive Dehn twists you gan get $\phi' = D_{c'}$. We've gone from the situation of having m Dehn twists, we have just one. If I had more time I'd try to justify this to you, but I don't.

How do I do this on the symplectic level? Attach more handles. So using this we can construct a symplectic manifold (X'', ω'') such that (X, ω) embeds and $\delta(X', \omega') = (M'', \xi'')$ with open book (Σ', ϕ'''), where $\phi''' = D_{c'}$ and $\delta X''$ is strongly convex.

Exercise 6 *M'' is the Euler class -1 circle bundle over Σ'' where Σ'' is just Σ' with a disk glued on.*

Exercise 7 *As time gets short, more and more exercises start appearing. If Y is the D^2 bundle voer Σ'' with Euler class $+1$. Then Y supports a symplectic form such that δY is strongly concave and δY will be $-M''$. Also show the induced contact structure is contactomorphic to ξ''.*

So what do we have here? We've constructed (X'', ω'') with strongly convex boundary (M'', ξ'') and then (Y, ω) is strongly concave and the contact structures match up so you get your closed manifold. The cap is the handles and then (Y, ω).

As you can see, the idea is to fiddle around in the mapping class group to get the monodromy under control, so that you can cap off.

Let me remind you, we went through the basics of contact structures and foliations, perturbing foliations to contact structures and symplectic handle attachments. I hope I've given you some idea of some of the rich interplay among different parts of mathematics that occurs in contact geometry.

[Can you tell us how Eliashberg's proof differs?]

He attaches a handle transverse to the binding. It's delicate but gives different information.

2 Ozsvath

Okay, I would like to say something about knot Floer homology. I'd like to say more, but unfortunately this is the amount I'll say. Everything I say is joint work with Zoltán Szabó. Some of this work is also [unintelligible]. This work is about understanding how to compute knot Floer homology along the lines, well, what is now familiar from Khovanov's theory.

So Khovanov tells us that really the skein exact sequence for the Alexander polynomial is what you want. You want a pair, my conventions will be random. For a positive crossing, you want a relation between an invariant for a singular link and its resolution. For a negative crossing you want it in the opposite order. I didn't define knot Floer homology but I assume everyone knows. So you could draw a cube of resolutions by iterating this, with edge maps.

I'd like to describe a candidate for what to put at a singular resolution. There are two candidates. Well, so be it, here's the candidate. Both of these were suggested by [unintelligible]. The candidates are the following. In the Heegaard diagram for a knot, one associates to a crossing a β circle going around like this and α circles going around the regions. So what should one associate with a crossing. You can use planar diagrams. If you have a crossing of this type, you replace each crossing with β-circles with other β circles inside of them, and α circles for edges. You place w and z basepoints. The w basepoints can't be crossed by disks, while the z can. I have as many z as I have edges. The differential will give the sum over all intersection points and homotopy classes of $U_1^{n_{z_1}} \dots U_n^{n_{z_n}}$. There's also a pair of basepoints. I should be a little bit honest. I'm going to draw all my knots in braid position. That edge will look different. We set that to be zero. Now this looks like it has more structure but it doesn't. You might think it's easier to compute because it's planar but the differential is pretty complicated.

Okay. So for the singular crossing you could put the figure eight. [Picture.] This is what you do when you resolve the crossing.

So what I'd like to say is that there is indeed an exact triangle with these digrams. The idea is that one has here a collection of immersed β circles. So we can just think of this as the invariant of a singular link. I would count only smooth holomorphic curves.

I can actually, um, um, right. There's an alternative description that has a more intuitive. There's an alternative candidate for what to put at the singular point that has a more intuitive definition.

So really, given how easy it is to prove exact triangles it's not that bad a thing to do. What one needs to understand is, one needs canonical generators. To that end we have to compute the Floer homology, let's try to understand the homology of this picture. Here we have four intersection points. We need to count flow lines. What you see is that the homology groups here are, we have the x_1y_1 gets mapped to x_1y_2 with differential $U_3 - U_4$ and $x_2y_1 \to x_2y_2$ via the same differential.

In fact, what we'd like to do, we define the maps to be tensoring with one of these generators. The problem is that you need pairs of cancelling triangles. What we need, the double composite to be chain homotopic to zero, we need the associativity of triangles. I have a generator coming in over here, and what we want to see is that the pairings give zero, like there are pairs of canceling triangles. We want the generators to be points connected by pairs of triangles. One of the points in my picture isn't closed.

[Pictures.]

Let me say it the way I like to, which is perhaps not the simpler way. We have a higher term and a lower term. Now $X_{\beta\gamma} \to Y_{\beta\gamma}$, we can look at the mapping cone, nad there's a differential down here which is multiplication by v. My point is that this isn't a cycle, but you can cook up a cycle from it. You can think of it as a differetntial graded algebra, [unintelligible]

So $\Theta_{\beta\gamma}$ will be $Y_e + X$.

Since we made the identification between two of the U, there's a verson that relates the three versions. [unintelligible]

We can compute the bottommost terms. The Floer homology of a completely singular link is actually lovely.

We want a chain complex over $\mathbb{Z}[U_1, \dots, U_e]$. We take our knot in braid position, singularize all our crossings, and get a module over this polynomial algebra supported in a single degree. In that fixed degree it is $\mathscr{A}$, which should start to smell reminiscient. In order to get a nice looking theory I'd like to work over a Novikov ring. I need additional basepoints, which should give me extra t powers. It's not so interesting in this picture. So now we allow finitely many t^{-1} and infinitely many positive powers.

The algebra is the following: introduce one relation for each crossing. The relations should be U_aU_b and $U_aU_b = T^2U_cU_D$. We introduce the relation $\prod U_{out} = T^{\#vert} \prod U_{in}$.

[At this point I admitted to myself that I had been completely lost for maybe 45 minutes and stopped trying to take notes.]

3 Fintushel-Stern

Today's lecture is on smooth simply connected 4-manifolds with $b^+ = 1$. There's $\mathbb{CP}^2, \mathbb{CP}^2 \# k\overline{\mathbb{CP}}^2$, $S^2 \times S^2$, and then we have to work. Start with any smooth simply connected X and equip it with a Riemannian metric. Recall that on Monday, I said the moduli space of solutions to the Seiberg-Witten equations has good properties given that there are no reducible solutions. But that was a lie.

Okay you need to perturb the equations. This is accomplished by choosing a self-dual 2-form η. It's still the case that reducible solutions cause the problems, and they're still of the form (A, c), where A is a connection on the characteristic line bundle and c a section of an associated vector bundle. The equation they satisfy is $F_A^+ = i\eta$. Let's recall what some of this stuff means. Once you have a metric you have a $*$ operator, and $\Omega^\pm$ are the eigenspaces of the $*$ operator. We can split $H^2(X, \mathbb{R}) = H_g^+(X) \oplus H_g^-(X)$. Then $[F_A] = \frac{2\pi}{i} c_1 L$. Then $-2\pi c_1(L)^{+,g} = \eta \in H^+(X)$, so $\underbrace{(2\pi c_1(L) + \eta)}_{\in H_g^-}{}^{+,g} = 0$. So this is the same thing as $(2\pi c_1(L) + \eta) \cdot v = 0$ for all $\underbrace{v \in H_g^+(X)}_{\dim = b_X^+}$.

So bad solutions cut out a codimension b^+ affine subspace. If $b_X^+ > 0$ then $SW_{X_{g,\eta}}(L)$ makes sense. If $b_X^+ > 1$ then we can connect any two good points by a path of good points. So in that case, $SW_X(L)$ is well-defined, just choose one and calculate. If $b_X^+ = 1$ then in a path (g, η) you might run into reducible parameters (g', η').

So what we can see is that, really what is determining the value in any case, when $b_X^+ = 1$, then $H_g^+(X)$ is spanned by a single element. We're choosing an orientation of the line. We choose a period point, in H_g^+ whose square is one, so there are two choices and that determines an orientation. The value of $SW_{X_{g,\eta}}(L)$ depends on the sign of $(2\pi c_1(L) + \eta) \cdot \omega_g$. Then we write $\begin{cases} SW^+_{X,\omega_g}(k) = SW_{X_{g,\eta}}(k) \text{ if } (2\pi k + i\eta) \cdot \omega_g > 0 \\ SW^-_{X,\omega_g}(k) = SW_{X_{g,\eta}}(k) \text{ if } (2\pi k + i\eta) \cdot \omega_g < 0, \end{cases}$ There is a a wall-crossing formula that says the difference as you cross a wall, by ± 1.

Okay. So there is a small perturbation invariant $SW_{X,\omega_g} = \begin{cases} SW^+_{X,\omega_g}(k) \text{ if } k\omega_g > 0 \\ SW^-_{X,\omega_g}(k) = SW_{X_{g,\eta}}(k) \text{ if } k\omega_g < 0, \end{cases}$

This also has a wall-crossing formula.

Exercise 8 *Show that if $b^- \leq 9$ then SW_{X,ω_g} is independent of ω_g.*

As a hint, use well-crossing and the fact that $SW(k) = 0$ if $d(k) < 0$.

So $SW_{X,\omega}$ satisfies all the usual properties of SW for $b^+ > 1$. For example, if g is a positive scalar curvature metric for X then $SW_{X,\omega_g} = 0$. So the small perturbation Seiberg Witten invariant of $\mathbb{CP}^2 \# k\overline{\mathbb{CP}}^2$ for $k \leq 9$.

Let's do some examples. We'll start with the Dolgachev surfaces. We'll start with $E(1) = \mathbb{CP}^2 \# 9\overline{\mathbb{CP}}^2$. This fibers over S^2 with generic fiber the torus. We can look at $E(1)_p$, the multiplicity p log transform. This means we think of the fiber as $S^1 \times S^1$ and the neighborhood of it as $S^1 \times S^1 \times D^2$, and then perform a p-framed surgery. It's a good exercise to try to distinguish this manifold from $E(1)$ and then if you do, you've made a mistake because they're diffeomorphic. BUt $E(1)_{p,q}$ for p, q relatively prime, will give a manifold homeomorphic to $E(1)$ but not diffeomorphic, assuming here that $p, q \geq 2$. The first exotic four-manifold was $E(1)_{2,3}$, that gave a counterexample to the 4-dimensional h-cobordism theory. This was Donaldson (1985). These are called Dolgachev surfaces. You can calculate $SW_{E(1)_{2,3}} = t + t^{-1}$. The primitive class is the fiber divided by six, and that's represented by t. Remember that $E(1)$ has Seiberg-Witten invariant zero. So they're not diffeomorphic. This was much more difficult to do with Donaldson invariants.

I'm sure it's a good exercise to calculate that p, q classify these. That's a result of Freedman and Morgan.

Let's to knot surgery on this manifold, $E(1)_K$. Then let's look at K_n a twist knot, with $2n-1$ right handed half twists in a double of the unknot. Then $\Delta_{K_n}(t) = nt - (2n-1) + nt^{-1}$. Then $SW_{E(1)_{K_n}} = -nt + nt^{-1}$.

Did I write down that $SW^-_{E(1),T} = \sum_{n=0}^{\infty} t^{2n+1}$, evaluated on the ray given by multiples of the fiber. If you multiply throught you can make this calculation.

So right here alone there's an infinite family homeomorphic and not diffeomorphic to $E(1)$. This always lives through blowups. So it's unknown whether there are minmial four-manifolds homeomorphic but not diffeomorphic to $\mathbb{CP}^2 \# (9+k)\overline{\mathbb{CP}}^2$ for $k > 0$.

So for $b^- = 8$ there is the Barlow surface B constructed around 1989 by Kotschick, which is not homeomorphic but not diffeomorphic to $\mathbb{P}^2 \# 8\bar{\mathbb{P}}^2$. Then there's the Park manifold P (J. Park, 2003) which is homeomorphic but not diffeomorphic to $\mathbb{P}^2 \# 7\bar{\mathbb{P}}^2$. Then within a year, there was one with $b^- = 6$, which I should call SS for fairness, discovered by Stipsicz-Szabo.

There's a time-dependent thing. Let me say, there's a fundamental dichotomy, for all known four-manifolds, if the homeomorphism type of X contains more than one smooth structure, it contains infinitely many. This wasn't true, and what I'd like to describe to you is how to correct this and restore the dichotomy, which was done about two years ago by Ron Stern and myself.

Are there any where there's known to be only one?

[There's not a single four-manifold for which we can classify all the smooth structures.]

Okay, I want to take as a starting point $E(1)$, the elliptic fibration which contains singular fibers I_8, I_2, and $2I_1$. So I_n consists of a cycle of n 2-spheres with self-intersection number -2. So I_2 can be deformed into a double node. These two, the cycles that vanish to these are homologous.

Okay, so you have $E(1)$ looking like this. Let me draw the same picture I erased and had so much trouble telling you about. In the double node neighborhood, let's take a fiber and do a knot surgery on that fiber. Let's remember what that is. We replace a neighborhood of that fiber with $S^1 \times 0$ surgery on the knot. Look at the loop Γ around the clasp. It's a loop in $pt \times S^3 \backslash K_n$. Now push to pt', and isotope Γ off itself to Γ' with linking number one with Γ. When we glue this into $E(1)$, sending the longitude to the boundary of D^2, and we'll glue the meridian to the vanishing cycle of the double node. So we see that Γ bounds an annulus to Γ' and then a -1 disk. So Γ bounds a disk of linking number $+1-1-1=1$.

Now recall the effect of knot surgery on a section. So back in the nodal fiber we've introduced some genus. We have a genus one pseudosection of self-intersection number -1. We've glued in a Seifert surface with Γ lying on top of it, generating homology but trivial in the ambient manifold. If this were self-intersection zero, you'd do surgery on it and turn the torus into a sphere, but here instead you get an immersed sphere, and you get a nodal pseudosection. You blow up once and get rid of the double point, but you have to add four to the self intersection number of S. If you wanted manifolds with $b^- = 8$ and $b^+ = 1$, well, you have -5 and -2 transverse to one another. That's something we can rationally blow down. So we know the Seiberg Witten invariants of this and of the blowdown.

Let me just show you how to see $b_- = 7$, and then all the ideas will be there. Blow up a nodal fiber, add to the section, and resolve the double point by replacing the crossing with the annulus. In the -8-singularity, you can see a lot of -2s to play with. So you can blow down the $-7, -2, -2, -2$. So if we use K_n, let's call the blowdown X_n. Then SW_{X_n} will have exactly two classes, and both will have Seiberg Witten invariant $\pm n$. All the others will give 0 or ± 1. So this gives an infinite sequence of manifolds equivalent to $\mathbb{CP}^2 \# k\overline{\mathbb{CP}}^2$ for $8, 7, 6$. A few moments later Park-Stipsicz-Szabo found the same for $b^- = 5$. So the next question is b^- is four or less.

If it's not true that $\mathbb{CP}^2$ has exceptional structures, we should figure out why, what the excuse is.

I'd like to thank you for your patience and interest. It's been a great experience except for the lack of sleep.

[How do you construct the elliptic fibration?]

You can use words in the mapping class group of the torus, and [*unintelligible*] classifies all possible singular fibers. Or find a good pencil in $\mathbb{CP}^2$ to blow up, and just be careful to see what the possible exceptional fibers look like.

4 Morgan

We're in the midst of discussing a theorem, I'll put that on the board.

Theorem 4 *(Perelman)*
If (M, g) is a finite time Ricci flow on M a compact 3-manifold, and we have a sequence (x_n, t_n) in $M \times [0, T)$ which is a blowup sequence, meaning $Q_n = R(x_n, t_n) \to \infty$ as $n \to \infty$. Then after passing to a subsequence there is a limiting flow for the rescaled and time-shifted $(M, Q_n g(Q_n^{-1}(t - t_n)), (x_n, 0))$. Here this is an ancient solution, existing for all negative

time. It's not flat or negatively curved. Each time slice (M_∞, g_∞) is complete and of bounded curvature, and lastly the limit is κ-non-collapsed for some κ depending only on $(T, (M, g(0)))$.

You have to control injectivity radius $inj_{(M,g_n(0)}$ at x_n, bounded below by a constant independent of n, of the rescaled flow, and get bounded curvature at bounded distance from the basepoint (again uniformly independent of n). This is because the higher derivatives of the curvature are controlled by curvature bounds, so that will be taken care of.

Last time I finished up with the $\mathscr{L}$-function, used it to show that the sequence is κ-non-collapsed at every point for a universal κ. So that addresses the injectivity radius at the basepoint. This will come from the bounded curvature at bounded distance along with κ-non-collapse. This will pass to the limit because κ-non-collapse is geometric and passes to a limit.

I want to talk about dealing with bounded curvature at bounded distance what these κ-non-collapsed solutions look like.

If I were Bourbaki I would put a funny twist like an S in the margin to indicate a funny twist in the argument.

We're going to have to do a funny induction on space and time, assuming what the limits will look like. I can't just prove the bounded curvature bit without knowing what the solutions will look like.

The structure of κ-solutions, meaning, something that has all the properties of the solution $M_\infty, g_\infty(t)$ in the theorem.

How do you understand what these look like? You use the $\mathscr{L}$-function again.

Go back to any time slice τ. I'm using things parametrized by negative time. τ will go from 0 to ∞, being $-t$. I look at the ℓ-geodesics back to the time slice $M \times \{-\tau\}$. I'm going to need the lemma again that says there's a a q_τ with a short geodesic γ to q_τ so that $\ell(\gamma) \leq 3/2$, where

$$\ell(\gamma) = \frac{1}{2\sqrt{\tau}}\mathscr{L}(\gamma) = \frac{1}{2\sqrt{\tau}}\int_0^\tau \sqrt{\tau'}(R(\gamma(\tau')) + |X(\tau')|^2)d\tau'$$

This is not a geodesic in this whole space, and ℓ is not a length function. So I produce q_n in time τ_n with short ℓ-geodesics.

Now it turns out that again properties of the length function tells me that [unintelligible]. So now consider $(M_\infty, \tau_n^{-1} g_\infty(\tau_n t), q_n - 1))$. This is a blowdown limit rather than a blowdown limit.

I think of this as a flow defined from $-\infty$ to -1. What used to be time slice $-\tau$ gets shrunk to -1. This is a sequence of Ricci flows where I throw away what is happening above q_n. So these converge, after passing to a subsequence, to a limit $(M'_\infty, g'_\infty(t), (q_\infty - 1))$, for $-\infty \leq t \leq -1$. But you get a very special kind of thing in the limit, a gradient shrinking soliton.

There's something called the Harnak inequality telling you that the curvature at different points and times are related by an exponential depending on distance in space and time.

As a family of Riemannian manifolds, they're all the same up to a scaling factor. So we have to start by understanding what the gradient shrinking solitons look like. So what are these in dimension three? This is something we can classify. The universal cover of one is either $(S^3, -tg_0)$ or $(S^2 \times \mathbb{R}, -tg_0 \times ds^2)$. So we can classify all of them as either compact round shrinkers, like the dodecahedral space or a lens space. We could take $\mathbb{RP}^2 \times \mathbb{R}$ or the twisted line bundle over $\mathbb{RP}^2$. If you have $S^2 \times S^1$, that solution wouldn't be κ-non-collapsed. So we won't be able to deal with $\mathbb{RP}^2 \times \mathbb{R}$, we'll assume that we don't have anything like that. You could think, if you want, only about orientable manifolds.

This M'_∞ and associated information is called the asymptotic gradient shrinking soliton of the κ-solution.

So how much can we learn about κ-solutions. So suppose we have a κ-solution whose asymptotic gradient shrinking soliton is compact. Well, it's a general fact about limits, if the limit of a geometric sequence is compact, then all of the things in the sequence are eventually diffeomorphic to the limit. So M is compact and diffeomorphic to the gradient shrinking soliton. So things are getting closer and closer to round. We have Hamilton's theorem that round things get rounder and then shrink away in finite time, and if you get close enough you only get rounder. How close you are to round in invariant. As I flow forward, then, it gets only closer to round. Se the whole κ-solution is just a round manifold shrinking. So when the asymptotic gradient shrinking soliton is compact, the M actually is its gradient shrinking soliton.

We want to understand singularity development. To do that you have to study κ-solutions. Then to understand those we need to understand the gradient shrinking solitons. The S on the board means that the flow of the logic isn't so simple.

So what happns if the gradient shrinking soliton is $S^2 \times \mathbb{R}$. Suppose that M^3 is a compact Riemannian manifold with nonnegative curvature, containing a long almost-cylinder. In it, we can find for appropriate ϵ, a region very close to $S^2 \times (-\epsilon^{-1}, \epsilon)$. The manifold somehow completes, and this is a metric statement, the metric is close to this product metric. Then what can you say about the topology of M?

Let's make life even easier. First let's assume the curvature is positive. Then this manifold admits a round metric and is covered by the three-sphere. But what manifolds allow this kind of thing? Either S^3 or $\mathbb{RP}^3$. One caps off with two balls, the other with a nontrivial $\mathbb{R}$-bundle over $\mathbb{RP}^2$.

In the noncompact case, then M is $\mathbb{R}^3$, a classical theorem of differential geometry, and it has one positively curved cap.

Now let's talk about if the Riemannian curvature is greater than or equal to zero. Here I'll use the strong maximum principle. This tells me that either the curvature is strictly positive or it locally splits as a surface cross a line.

In the world of κ-non-collapsed, you can either have $S^2 \times \mathbb{R}$ or an $\mathbb{R}$ bundle over $\mathbb{RP}^2$. Then you have the more interesting ones, S^3 and this one-capped cylinder. In the cigar you get large curvature and low injectivity radius, so that's different.

Understanding gradient shrinking solitons has helped us to understand some of what the κ-solutions look like. The result is a focus on what are called canonical neighborhoods. You need to prove a compactness condition for κ-solutions. So any sequence of κ-solutions with basepoint [unintelligible], everything will converge to another κ-solution.

So you can prove the following:

Every point in a κ-solution has a neighborhood of one of the following types:

1. compact round
2. center of an ϵ-neck.
3. in the core of a cap, either a three-ball or a punctured $\mathbb{RP}^3$ with positive curvature.
4. manifolds of positive curvature, compact, but not round, either S^3 or $\mathbb{RP}^3$, and completely bounded geometry.

This comes out when you flow in one of these round ones. If we let this solution go far enough, it will become arbitrarily small. So it should go from having something horrible at $-\infty$, with a long neck, the neck should collapse down eventually since it has to shrink to nothing in finite time.

It's not known how many of these things there are. These are not easy to come by, you can make one example.

So this leads to a notion of a canonical neighborhood. If you have a neighborhood close to one of these, you call that the canonical neighborhood. You might put an ϵ on that controls how close your metric is to the product, how long your neck is, and so on.

That's our discursion into what κ-solutions look like, and we ended up with these models of canonical neighborhoods. Now I'm ready to state a theorem, let's call it a lemma.

Lemma 1 *Suppose we have a Ricci flow $(M, g(t))$ up to time T. Suppose there are constants $Q_0 < \infty \ldots,$*

[Very long pause.]

I'm just trying to think [unintelligible]formula [unintelligible]

Suppose we have a sequence of Ricci flows $(M_n, g_n(t))$ for time up to T_n which are bounded by some $T < \infty$, and that $(x_n, t_n) \in M_n \times [0, T_n]$ $Q_n = R(x_n, t_n) \to \infty$ and suppose that all (y, t) with either $t < t_n$ or $t = t_n$ with $R(y, t) \geq 2R(x_n, t_n)$ have ϵ-canonical neighborhoods. Then so does (x_n, t_n) for all n sufficiently large.

We have these at zero, because there won't be points of high curvature near zero. I should have said $(M_n, g_n(0))$ normalized, with Riemannian curvature at most 1 and the volume of a ball of radius 1 at least half the volume of a Euclidean ball.

So here's the idea of the proof. Take a blow up sequence based at the (x_n, t_n). We want to form a blow up limit. Since the original manifolds are [unintelligible], all solutions are κ-non-collapsed. We look for points (y_n, t_n) where the curvature is arbitrarily large compared to (x_n, t_n), but the rescaled distance $\sqrt{Q_n} d_{g_n(t_n)}(x_n, y_n)$ stays bounded by, say, C.

Because of the canonical neighborhoods you can take a limit out to the first such point. First you have to see that this distance stays bounded below. Then you take an incomplete geometric limit out to the first bad point, from the compactness and differential equations satisfied by the canonical neighborhoods.

It turns out that you look at the neighborhood of the singularity, and you get little ϵ-necks near the singularity. In fact, this is positively curved. So then you can take a limit around the missing point, the y_∞, and you get a cone. So you get a Ricci flow that ends in the open piece of a non-flat cone. But a cone has a direction of trivial sectional curvature. It ends up in a non-flat cone, but it has a flat direction, so is a metric product by Hamilton's maximum principle, but a cone is not a metric product, so this sequence doesn't exist. So you have a κ-solution and so there's a canonical neighborhood out near the limit, meaning that eventually there's one near the x_n for high enough x_n.

I'd better stop, I can see the goons in the back.