Low Dimensional Topology Notes
July 14, 2006

Gabriel C. Drummond-Cole

July 14, 2006

1 Morgan

[My talk this afternoon was cancelled. I felt like I couldn't finish this morning. I felt like I couldn't finish, so I postponed my final talk to 4:30.]

What I showed you last time, pushed a little further, will allow you to prove the following.

Theorem 1 *For every $T < \infty$, there exist $r > 0, \kappa > 0, r_0 > 0$ such that the following holds: for any normalized initial conditions $(M, g(0))$ (meaning the volume of a radius one ball is bounded below by half the Euclidean volume, and $|Rm| \leq 1$. and any Ricci flow $(M, g(t))$ for $0 \leq t \leq T \leq T_0$,*

1. *The flow is κ-non-collapsed on scales $\leq r_0$.*
2. *Any point (k, t) with $R(x, t) > r^{-2}$ has a canonical neighborhood*

The κ condition was a bound on the volume of a ball. Then scales of r_0 means that you only have the bounds for $r < r_0$.

So consider $(M, g(t))$, and assume the flow doesn't extend to T.

So let's look at $(M, g(t))$ for t close to T. What is happening as the singularity develops? The curvature has to be blowing up, from Hamilton, so we look at the regions of large curvature, all of which have canonical neighborhoods. if "large" means the hypothesis of the theorem, then we have canonical neighborhoods all over M_{large}. We have a list of four canonical neighborhoods

1. Compact round
2. S^3 or $\mathbb{RP}^3$ of bounded geometry

3. necks, metricly $S^2 \times (-\epsilon^{-1}, \epsilon^{-1})$
4. a cap B^3 or $\mathbb{RP}^3$.

We know the metric in either of the first two cases is standard throughout the component. So we know that the neighborhoods are all necks or cores of caps.

Look at a component of M_{large}. Every point looks like it's in an ϵ-neck or in the core of a cap.

Lemma 1 *Every component of this M_{large} is contained either in a tube or a tube with one or two caps on the ends.*

A tube is just a longer version of a neck. It's not metricly standard because as you move down the tube you might have the radius changing slightly. Nevertheless, as you move down the ϵ-tube, passing between ϵ-necks, things almost line up and you get a topological product. If you cap off both ends, you get $S^3, \mathbb{RP}^3$, or $\mathbb{RP}^3$, all of which we understand. If you cap off only one end you get either a disk or a punctured $\mathbb{RP}^3$.

Now let's go to the singular time and see what's going on.

Define an open subset $\Omega \subset M$ where $x \in \Omega$ if $\limsup_{t \to T^-} R(x,t) < \infty$. Then using canonical neighborhoods we can see

- Ω is an open subset of M.
- $g(t)|_\Omega \stackrel{C^\infty}{\to} g(T)$ on Ω_T.
- $R : \Omega \to \mathbb{R}$ is bounded below and proper.
- For any connected component Ω^0 of Ω, every end is a horn $S^2 \times [0,1)$, a union of necks with curvature going to ∞.

For example, you might get a single simple two-sphere and get two horns, one on each side, as a tube crushes down at one slice. It could be a lot more complicated than this.

I'm going to divide the components of Ω into two types. Choose some $\rho < r$. Then $\Omega_\rho = \{x \in \Omega | \Omega(x) \leq \rho^{-2}\}$. So there are components of Ω containing a point of Ω_ρ, and then there are those which do not.

There are finitely many components Ω^0 of Ω containing points of Ω_ρ. Then Ω^0. The other components look like double horns, capped horn, or one of the closed components we've already talked about. We have a finite number of components with not -very-large curvature, and each one of those has a finite number of ϵ-horns. That's what these neighborhoods look like.

Now what are we doing with this picture of the singular metric? In other words, how do we perform surgery to turn this back into a compact manifold?

1. Throw away all components of Ω that do not meet Ω_ρ.
2. For a component Ω^0 of Ω meeting Ω_ρ, in each horn, we find a δ-neck for some δ we have to determine, and cut off the horn at the central S^2 of the neck.

The main invariant of a neck is the scale, which goes to zero as the curvature goes to ∞. But the further you go down the tube, the better control you get about how close the metric is to standard. So you find a δ-neck, doing surgery with δ which depends on time, and is not for the fixed ϵ. So we cut the horn off and throw away the noncompact pieces. A neighborhood of these boundaries loks like $S^2 \times [0, \delta^{-1}]$ rescaled to look very small. So I got a finite number of these components that I cut the horns off of and all the other components that I threw away.

On the subset I constructed there's a smooth metric $g(T)$. Now I have to figure out how to complete these ends to get back to a closed manifold to start up a Ricci flow again. Now we have to choose what to glue in. The easiest thing, maybe, is to define a metric on $\mathbb{R}^3$, define g to be $SO(3)$-invariant, and metricly the end of $\mathbb{R}^3$ will look lie $S^2 \times [0, \infty)$ and then the core will be a 3-ball with positive curvature.

This is called the standard initial condition, and then you form the Ricci flow on this, $\mathbb{R}^3, g(t))$. You have to prove things about this Ricci flow.

1. It exists uniquely, which we don't have in general. Imagine you have a flat metric on $\mathbb{R}^3$, then you could flow with a constant flow. But you can put that in S^3 with some positive curvature and then the flow on the sphere will give positive curvature. Then you can restrict.

   I focus on this for a psychological reason, this is the only place we know Perelman made a mistake. He wrote down an ODE to show uniqueness, and his ODE had a singular point at the origin. But he know how to fix it right away.
2. The curvature is positive for $t > 0$
3. The flow exists on the interval $[0, 1)$. At time 1 the metric completely degenerates in the sense that the distance between any two points goes to zero as t goes to 1. The positive curvature is like a Pacman, it eats the whole cylinder, which at the whole time is shrinking down trying to get to be a line, but Pacman is faster than convergence to a line, so he eats it all before it becomes a line.
4. For any $[0, 1-\theta]$, up to but not including one, there exists compact X on $\mathbb{R}^3$ such that the flow $(\mathbb{R}^3 - X, g(t)|_{\mathbb{R}^3 - X})$ on the complement is close to a [unintelligible]cylinder.

I can rescale this cap until the curvature agrees and I have to interpolate between the metrics using a partition of unity.

Topologically. we've thrown away the compact components, and the tubes, in general, I have two horns I've concentrated on, and I've cut this and capped it off. That's an ordinary surgery on the core two-sphere of the annular region. So in the capped case, I have thrown away the cap and recapped with a ball. So I've done nothing or cut off an $\mathbb{RP}^3$ summand. So the topological effect of surgery is three-fold.

1. In the necks, you do ordinary 2-sphere surgery.
2. In the necks you remove B^3 or $\mathbb{RP}^3$ and replace with B^3.
3. We remove standard components, either compact positively curved, or S^2 bundles over S^1.

What we want to do is to take the new compact Riemannian manifold and use it as the starting point for the new Ricci flow. It's important not to view this as a sequence of disconnected steps. We're going to glue this together to create the topology, but since there's a change in topology it will not be a smooth manifold.

Let me draw a model for four-dimensional space-time.

Now let me draw a more accurate picture of what would happen on one of these necks. At the surgery time I replace the cylinder with two caps. You have these tubes which are removed and replaced with caps. So the flow will go up but not down at these times. So you have manifold with boundary points and then singular points at the end of the boundary. So this is $\mathscr{M}$ which comes with $t : \mathscr{M} \to \mathbb{R}$ and then $\underline{t}^{-1}(t)$, the slices, are compact three-manifolds M_t.

Exercise 1 *If the postoperative patient satisfies the Geometrization conjecture, then so does the pre-op.*

The main result out of Ricci surgery is that we need to show that this operation, this surgery, can be done repeatedly to construct a Ricci flow with surgery defined for all time. The surgeries might happen infinitely often in finite time. So you show that in any compact time interval there are only finitely many surgeries. As far as we know you can go all the way to ∞ and get infinitely many surgeries.

This afternoon I'll try to indicate to you the problems in repeating this process. I have to have some length function for κ-non-collapsed, and then to get canonical neighborhoods you do a blow up operation.

For the length functions, you can't always push backward, and blow up won't give you an ancient solution. But you can iteratively prove both of these, and I'll do that in the first half of the afternoon, and then I'll finish off to tell you how to get these two conjectures.

2 Fintushel-Stern

I think it's been a great effort on your part to attend all of these. Even though you may not understand but two or three percent, but multiply that by 80 lectures a year and then thirty years and you're pretty sure to have a hit.

Remember how I said, it's hard to tell a story when you don't know the end of the story. I want today to summarize and then tell you a bunch of open problems.

So we starte with constructions of 4-manifolds. We figured out some topological invariants, c, χ, and t. Here c and χ are just different ways of describing σ and e, and then t is a $\mathbb{Z}_2$ which tells whether the manifold is spin (even).

We learned a little about the Seiberg-Witten invariant, which is a Laurent polynomial. $SW_X \in \mathbb{Z}[t, t^{-1}]$, where $t \in \mathbb{Z}[H_2(X, \mathbb{Z})]$. Let me again remind you of the constructions we came up with to get new manifolds.

1. The generalized logarithmic transform started with a torus $T^2 \hookrightarrow X$ with a neighborhood $N(T) = T^2 \times D^2$ and then $X^4/\backslash N(T) \cup_\theta T^2 \times D^2 = X_\theta$ where $\theta : T^2 \times \delta D^2 \to \delta(X \backslash N(T))$ and this is characterized by three integers $X_\theta = X_T(p, q, r)$. Now $(c(X_T(p, q, r)), \chi(X_T(p, q, r)), t(X_T(p, q, r))) = (c(X), \chi(X), t')$, where t' is 0 if and only if $t(X)$ is 0 and r is odd.

   a. So we're wisely choosing particular operations that we know are effective; if the torus was inessential, $\pi_1(X \backslash T) = 0$, then $SW_{X_T(p,q,r)} = SW_X \frac{t^r - t^{-r}}{t - t^{-1}} = SW_X(t^{r-1} + \ldots + t^{1-r})$, where t corresponds to T in $H_2(X)$. This assumes there is a vanishing cycle, so T sits inside of a nodal neighborhood. A philosophical statement is that we've never really used all of our variables, so this is just saying you can get rid of one of the variables.

   b. When T was nullhomologous, there was the Morgan-Mrowka-Szabo formula, which was very effective, which described the Seiberg Witten invariants of the p, q, r log transform in terms of the $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$ log transforms. This gave us knot surgery. If $X_K = X \backslash N(T) \cup S^1 \times S^3 \backslash K$ then $SW_{X_K} = SW_X \Delta_K(t)$.

      Then we could use Dehn surgery to unknot the knot, and the Morgan-Mrowka-Szabo formula gave us a way to understand this.

2. The other technique was rational blowdown. That's fresher in your mind so I won't describe it again.

3. blowup and blowdown, $X \# \overline{\mathbb{CP}}^2$.

4. generalized fiber sum, where you have $\Sigma_g \hookrightarrow X_1$, $\Sigma'_g \hookrightarrow X_2$, and $\Sigma_g^2 = -\Sigma'_g$, $X_1 \#_{\Sigma_g = \Sigma'_g} X_2$.

There's an operation I didn't mention, and that's to change orientation. So there's a question, does every homeomorphism type have a canonical orientation? A symplectic or complex manifold has a canonical orientation. For a symplectic manifold you posit $\omega^2 > 0$. The first question that comes to mind, does an arbitrary smooth manifold have a canonical orientation? Given X, is $\pm X$ homeomorphic to a symplectic manifold? Strangely enough, we don't know if this is true or false. This sounds like a stupid thing to be concerned about, but being left-handed I'm concerned about orientation. All of these are irreducible.

It's known thanks to the work of Taubes that if I have an irreducible symplectic manifold that $c(X) \geq 0$. That translates to the fact that regardless of orientation $c < 48/5\chi$ for a symplectic manifold, and the 11/8 conjecture and in fact the 3/2 conjecture, that is, $\frac{b_2}{|\sigma|} \geq 3/2$. I apologize for dwelling on changing orientation. Now in talking about the geography question, I'm going to suppose there's a canonical orientation and restrict to symplectic manifolds. We certainly have complex Kähler manifolds, here's 9χ and here's $2\chi - 6$. Every complex surface lies in this region or on the χ axis. So inside $E(n)$ there are lots of -2 curves, at least $n-3$ curves connected to a -2 curve. This is the configuration you can blow down. So that lets you fill out this whole region with symplectic manifolds. When you look at the number of Seiberg-Witten basic classes. In the computation we did, $SW_{E(n)} = (t - t^{-1})^{n-2}$, so there are $n-2$ basic classes. In the construction of manifolds in the region on the lower right, they have a lot of basic classes. This is a construction type problem, to find a counterexample, do there exist X with $c(X) < \chi - 3$ with fewer than $\chi - c - 2$ basic classes? Again, this is a fact, we construct manifolds with this many basic classes, are there fewer? Understanding this will help us understand [unintelligible]. You'd need to be a complex geometer to prove it, sort of a Noether formula for symplectic manifolds. There are other obvious questions. Are there any symplectic manifolds above the 9χ line? Are there any at all above the $48/5\chi$ line?

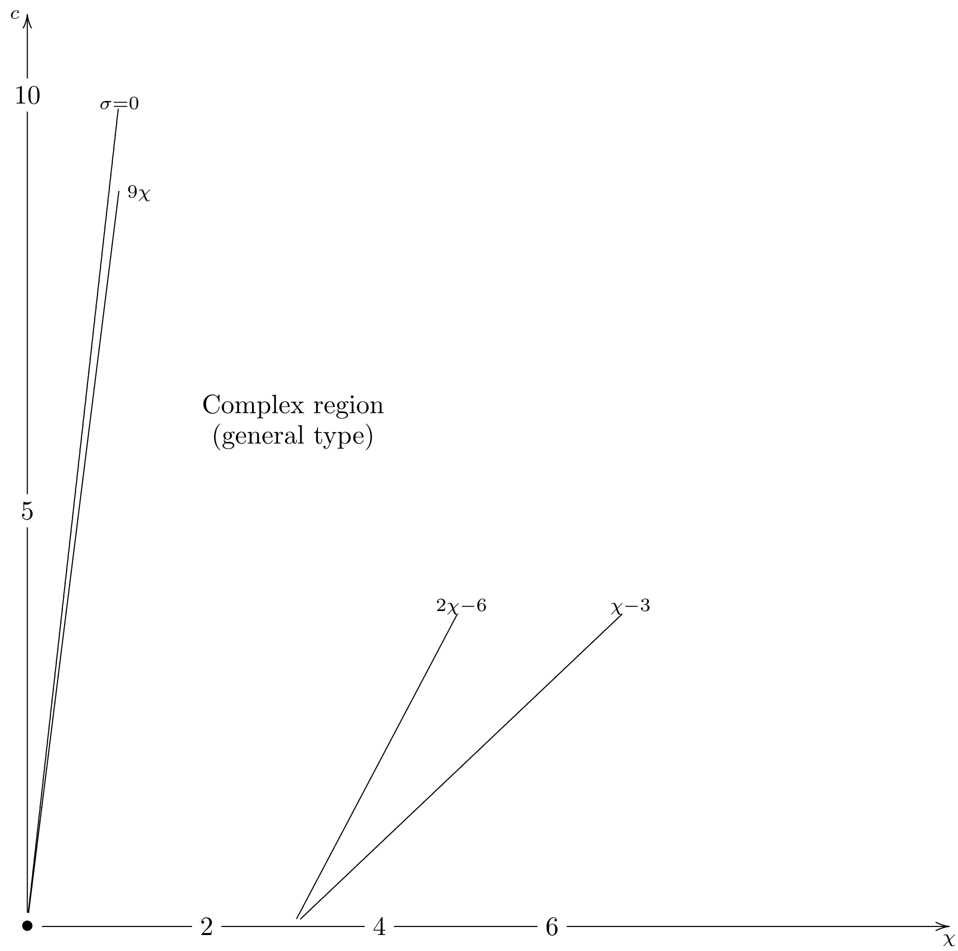

Okay, suppose I have two homeomorphic irreducible manifolds of general type. How do I pass from one to the other? The reasonable conjecture is that it's the list of operations we made earlier, since those are the only things we know how to do.

Now we know that there are only finitely many deformation types for a given homeomorphism type of complex manifolds (although we can usually get infinitely many smooth structures. As an aside, does every X^4 contain essential tori? If they have different deformation types, are they diffeomorphic?) Well, let me give you the first example of manifolds for which this is not known. These are the Horokawa surfaces lying on the Noether line $c = 2\chi - 6$, which have two deformation types. I would like to see that these can be passed one to the other, by these operations, to test the conjecture that these are the operations we need.

We'll construct them with branched covers. I gave you a formula for figuring out the c and

χ of a branched cover. I have, say six S^2 that run vertically and $2m$ that run horizontally. This is a bunch of intersecting lines, not a curve. You can observe what the singularities are where these lines intersect. I take a two-fold cover, and an exercise is it's the cone on the cotangent bundle of S^2. I end up with $H(2m-1)$, and the singularities are a cone on $\mathbb{RP}^3$, which I can resolve by replacing the cone by a -2 curve, a sphere of square -2. What does that mean? I look at a sphere with normal bundle with Euler class -2 and glue it in. This is the beginning of the resolution of singularities. I could also find a nice curve in $S^2 \times S^2$ with bidegree $6, 2m$. One way or another, I get a surface, and the exercise is that $c(H(2m-1)) = 4m - 8$ and $\chi = 2m - 1$. This is the way that Horikawa constructed his one family, so $H(2m-1)$ is a two-fold cover of $(6, 2m)$.

The other family is the following. Notice I said there are only finitely many deformation types in the general region. Now $S^2 \times S^2$ has infinitely many such types. Consider the Hirzebruch surfaces F_{2n}, topologically the double of the sphere of self-intersection number $2n$. Then the double has $-2n$. If this were odd I'd get $\mathbb{CP}^2$ and $\overline{\mathbb{CP}}^2$. So this, anyway, is $S^2 \times S^2$. So we have the two spheres, which we will call $S^{\pm}$ according to the sign of their self-intersections. So let's take 5 S^+ and 1 S^-. This is a disconnected branch set. You can represent that by a branched curve. We can take the two-fold branched cover of this. Call this $H'(4n-1)$. The exercise is that this is a manifold with $\chi = 4n - 1$ and lying on the Noether line. So $H'(4n-1)$ is the two-fold cover of $(5, 1)$ in F_{2n}. The hard part is not the numbers but that these are complex surfaces. The even harder part is:

Theorem 2 *$H'(4n-1)$ and $H(4n-1)$ are deformation inequivalent. These are the only complex surfaces along the line $c = 2\chi - 6$.*

Are they diffeomorphic? Here's some motivation that they're not. Let's look at $n = 1$. That's $H(3)$ and $H'(3)$. Well, a relatively easy exercise is that $H(3) \overset{\text{diff}}{\cong} E(3)$. But $H'(3)$ is $5 \times S^+$ and S^-, where these have ± 2 self intersection numbers. We take the twofold branched cover over this. What happens when I look at the lift? The sphere of square -2 becomes a sphere of square -1. So $H'(3)$ is not minimal.

Okay, and you can see that $H'(4n-1)$ is spin if n is even, while $H(2n-1)$ is always odd. So if n is odd, $n > 1$, is $H'(4n-1)$ diffeomorphic to $H(4n-1)$? This is an open problem. What I want to conclude with is that, can you pass from one to the other via the basic operations on the board? Well, $H'(4n-1)$ is obtained from $H(4n-1)$ via a ± 1 log transform on an essential torus.

So I'm not changing the Seiberg Witten invariants since it's essential. This is interesting because it's the first torus we've seen without a vanishing cycle.

Okay, two minutes, how do you prove this observation? One has a disconnected branch locus, and the other has a connected branch locus. I can braid them together with a half twist to make it connected. So one way to connect these is to look locally and get S^1 times this picture of connect sum with a twist. So upstairs this can correspond to a log transform. This braiding operation is worthy of graduate students to further think about.

The end of the story might be in another direction. We committed ourselves to this direction and this is where we're at. We understand something about existence. We also only know something about the number of such manifolds at a point. We conjecture that the list we made gives all the ways of passing from one thing to another at a point in this geography. This is where we're at. In some sense, this is the end of this particular process. We were talking last night, at the ending dinner in twenty years, what will they be talking about in topology. Hopefully it will be this, probably it won't.

[Could there be a symplectic form on both a manifold and its opposite?]

That's open.

[Is there any hope of new invariants?]

A lot of us have been dragged around by new invariants from the physicists. This other question is, does every manifold have simple type, either from Seiberg-Witten or Donaldson type? So have we milked everything from Seiberg-Witten? Donaldson theory is very sensitive to fundamental group; Seiberg-Witten invariants are insensitive. So there's a marriage there of opposites. Sensitive, or not, and somehow the same. Many people have thought in the Donaldson setting about higher gauge groups. Basically the $SU(n)$ theory is determined by the $SU(2)$ theory.

How many people here really are four-manifold graduate students? Six? I hope you have learned something here. I hope the rest of you got a few buzzwords, so you can nod your head, yeah, yeah, yeah, that's a very important property of mathematicians.

3 Morgan

As I was saying this morning, you needed these two basic results from Ricci flow to movo to Ricci flow with surgery, the κ-non-collapsed condition and the canonical neighborhoods.

Let me remind you of what happens with κ-non-collapsing in normal Ricci flow. I followed points back by backwards geodesics, and those had length at most 3/2, and I got a subset at time 0 of a reasonably sized ordinary volume, and then the integrand $e^{-\ell} \geq \delta$ and $vol \geq C$, and then we used the result to push a subset up along geodesics only making the volume go up. Then I made a local argument.

It's pretty clear what the problem is. What about ℓ-geodesics that run into the surgery caps and just stop. You needed ℓ-geodesics defined everywhere to apply the maximum principle. I can't do that in the naive case. To hit a point way over near the cut, you'd have to use something piecewise. But you can do the minimum argument if you know that geodesics back to the surgery caps have length at least $3/2 + \epsilon$

For the canonical neighborhoods you had $Q_n = R(x_n)$, with a rescaled metric, and you wanted to be able to flow these back and get an ancient solution. But there might be a cap in the way.

Okay. So the solution to both of these is the same. Let δ be the surgery control parameter. Then as $\delta \to 0$, the $R(y) \to \infty$, so the radius goes to 0 as δ does. I can say how tightly I'll control the metric on the surgery region where I do the cutting.

That turns out to be the freedom that allows me to arrange that the geodesics are long and that I get canonical neighborhoods.

So I define $\delta(t) : [0, \infty) \to [0, \infty)$. When we do surgery at time t, we use $\delta(t)$ as the control parameter. We can specify it a priori if we want. The scale at which we do surgery, the function will tell you the control parameter. So, δ of course always has to be positive, it's convenient for it to be decreasing, and it has to fall off sufficiently fast.

The advantage of making δ very small is, let me draw it big. Here's half of a δ-neck. This end goes off to the horn. Now I glue in my standard cap glued in to size δ. These converge to the standard solution. Now, it's true, the amount of time I have to control is only one unit in this scale. This is exactly what you need. If when you rsecale at x_n you always find a surgery picture, you know that you can find it in the flow of a standard solution.

That's exactly how you solve the canonical coordinates problem. So that gives me the canonical neighborhood not coming from a κ-solution but from the evolution of the caps. Those have the same compactness properties, and you'll still be able to make things work. Now this diagram also solves the κ-non-collapsed problem. Coming in from the side gives you length; coming in from the top the curvature is huge.

That's briefly the argument for extending these arguments from Ricci flow to Ricci flow with surgery. Now let me state for you the theorem that one proves. You might wonder what kind of function δ should be.

Theorem 3 *There exists an $\epsilon > 0$ and a decreasing sequence of positive constants $\kappa_0 \geq \kappa_1 \geq \ldots 0, r_0 \geq r_1 \geq \ldots 0, \delta_0 \geq \delta_1 \geq \ldots 0$, such that for any normalized initial conditions $(M, g(0))$, there exists a Ricci flow with surgery defined for all time, where the region $[0, \infty)$ is divided into the regions $[0, \epsilon), [\epsilon, 2\epsilon), [2\epsilon, 2^2\epsilon), \ldots,$ which I will call the intervals $0, 1, \ldots,$ where the surgery parameter is $\delta(t) = \delta_i$ in the ith interval, and the Ricci flow with surgery in $[0, 2^i\epsilon)$ is κ_i-non-collapsed on scales at most ϵ. and has canonical neighborhoods at points x where $R(x) \geq r_i^{-2}$.*

The constants must be chosen very carefully and inductively and by contradiction. So they're falling off pretty fast. We don't have any explicit estimates.

That's the actual theorem about the existence of Ricci flow with surgery for all time. ϵ goes with something like, gluing things and needing them to be isotopic. That makes sure that the description I gave you of manifolds covered by ϵ-patches is concrete.

All right. What time is it? Now we want to switch gears and see what we can do now that we see that the solution exists for all time.

Let me talk first about geometrization. If for sufficiently large t, we can prove that the manifold satisfies geometrization, then it does at all time.

I omitted something. Every time you do surgery you cut off half a δ-neck, and there's a fixed amount of volume lost that depends on δ, but it's positive. Under Ricci flow, volume grows at a fixed exponential rate. Every time you do a δ-surgery you can only cut a certain amount away. Each surgery up to time T cuts out at least a certain amount since up to T, δ is bounded away from zero, so there can only be finitely many in a given time interval. Surgeries that throw away components, well, there are only finitely many components, and the other kind of surgery can only add one component, and there are only finitely many of those.

So Hamilton studied what happened at ∞ if the Ricci flow existed for all time. This divided into pieces that were almost hyperbolic with torus boundary and other parts that were collapsed. In Ricci flow with surgery, for all sufficiently large time it will divide into things where the curvature is negative, very close to zero, and then collapsed things with short loops. These fit together along tori, the limit tori of hyperbolic pieces are very small so they fit with the collapsed pieces.

Hamilton argued that the tori are incompressible. I'm not going to give it, but if you suppose it's compressible look at a minimal disk, and what happens to the area of the disk? Its area is going to go negative in a finite amount of time. If you go further out in time, a limiting argument tells you that you can't have any such compressing disks.

If you're interested in geometrization, if you get this picture you might as well stop because it's Haken. If you get the hyperbolic pieces, you can stop.

The theory of collapsed manifold was first studied by Cheeger and Gromov, and developed F-structures, which in dimension three says that given certain conditions, the collapsed pieces are either Seifert fibered surfaces with short fibers or tori fibering over a 1-manifold, with a short loop in the torus. These are in general nilpotent; in this dimension they're Abelian because it's in two dimensions. So Shioya-Yamaguchi have a result, very delicate and relying on unpublished work of Perelman, say the collapsed regions are graph manifolds. I'm not sure about that part. I think the Chinese may have an explicit argument.

The other way to go is to be modest and go for the Poincaré conjecture, or maybe that's too modest, and go for the $3D$-space forms.

Theorem 4 *Say that $\pi_1(M)$ is a free product of finite groups and cyclic free groups, then M is a product of $3D$ space-forms connect sum with S^2 bundles over S^1.*

Theorem 5 *For 3-manifolds of this type, $M_T = \emptyset$ for some $T \geq 1$.*

We start with this kind of manifold, and start doing the surgeries, which will eventually do the decomposition, and eventually those disappear. As each one disappears, it's either S^2 fibering over S^1 or has a round metric.

So, why does this happen? Well, it's really for the same reason I hinted at with the compressing two-disks. There are three steps.

1. After a finite time all of the two-sphere surgeries are on trivial two-spheres.
2. After a finite time all π_2 are trivial. So either you have a homotopy three-sphere as universal cover or you have a $K(\pi,1)$ as universal cover.
3.

So why are the 2-sphere surgeries trivial after finite time? Here's some topology, you'll be happy to know. This is called Grusko's theorem.

Theorem 6 *Suppose G is a free product of $G_1 * \cdots * G_K$. Say a free group F maps onto G, then $rk\ F \geq \sum rk\ G_i$. He really proves you can write F as $F_1 * \cdots * F_K$ so that F_i maps onto G_i. In particular you can't write a finitely generated group as an arbitrary free product.*

So any time you cut a manifold along a two-sphere, you get a nontrivial decomposition into a free product. So after a certain number of cuts, all S^2 surgeries are along homotopically trivial 2-spheres. Let's follow a component after surgery. It could disappear, keep flowing, and if I do a surgery, I get something homotopy equivalent to the manifold I start with, along with some homotopy three-spheres. So X' has the same π_2 as X. The only way to destroy a π_2 is to kill that component. Follow the one that's homotopy equivalent until it disappears. If there's one of these arbitrarily far out, we can find a component $X(T_0)$ at time T_0, and then at any time $t > T_0$ there's an $X(t)$ which flows from $X(T_0)$ during normal flow, and moves to something homotopy equivalent during surgery.

Notice I don't need to deal with this for Poincaré, I can just assume I started with something prime.

So $S^2 \xrightarrow{\varphi} X(t)$ for $\varphi \not\simeq *$, with minimal area of ϕ equal to $W_2(t)$ where W_2 is a continuous function of t with $\frac{dW_2(t)}{dt} \leq -4\pi - \frac{1}{2}R_{\min}(t)W_2(t)$, bounded in the sense of forward difference quotients. So I can bound $R_{\min}$ below by $\frac{6}{1+4t}$. So you get the equation $\frac{dW_2(t)}{dt} \leq -4\pi + (\frac{3}{1+4t})W_2(t)$. So if there's a solution when this is an equality, you will get negative volume in finite time. This only arose because I assumed an all time existence of this X, so that this actually gives me a bound on when X disappears.

Look at the minimal sphere. If I can show the particular inequality for the area of this sphere, it will bound the minimal sphere. The formula for the area, what is $fracddtA(\varphi(S^2))$? It's simply

$$\int_{\varphi(S^2)} \frac{1}{2}Tr\ (\frac{\partial g}{\partial t})|_\Sigma da = -\int_{\varphi(S^2)} Tr\ Ric(g)|_\Sigma da$$

where this trace is $Rm(e_1,e_2)+Rm(e_1,e_3)+Rm(e_2,e_3)+Rm(e_1,e_2) = 2K(e_1,e_2)+K(e_2,e_3)+K(e_1,e_3) = \frac{1}{2}R + K(e_1,e_2)$.

So this is

$$-\int_{\varphi(S^2)} \frac{1}{2}Rda - \int_{\varphi(S^2)} K(e_1,e_2)da$$

Now we use that φ is minimal and the Gauss-Codazzi equation which says $K_\Sigma = K(e_1, e_2) + \det(II)$, so we can rewrite this as

$$-\int_{\varphi(S^2)} \frac{1}{2} R da - \int_{\varphi(S^2)} K_\Sigma da + \int_{\varphi(S^2)} \det(II) da.$$

Since this is minimal the determinant is negative, so we forget it. Then we have the inequality

$$\leq \frac{-1}{2} R_{\min}(t) A(\varphi(S^2)) - 4\pi.$$

So the homotopy equivalence along the surgery can be chosen to be a distance-decreasing map, and so it's area decreasing and so get the appropriate area drop, or at least not increase. This shows that eventually π_2 is trivial.

So either it's covered by a homotopy sphere or it's acyclic. What if you have finite π_1? It might bifurcate into other components. I want to follow them all. I have a finite collection of components generated by this one. In terms similar to the previous argument you can get a bound during which these will disappear in finite time. So take a nontrivial element $\alpha \in \pi_3(X(T))$. This can be thought of as a one-parameter family of two-spheres; Perelman did something more naive, but on the face of it something more twisted. It's easier in the end to analyze.

Lemma 2 *Assuming $\pi_2(X(t))$ is trivial, then $\pi_3(X(t)) \cong \pi_2(\Lambda X(T), *)$.*

So we reinterpret α, so we have a 2-sphere of homotopically trivial loops. For each trivial loop, we associate something we call the area. It's simply the area of the minimal disk spanning c.

[Is this based?]

It doesn't matter. I need them to be homotopically trivial. If I have a family of loops, say a S^2 worth of loops Γ, then I define $A(\Gamma)$ to be the $\max_{c \in S^2} A(\Gamma(c))$. I can call $W(\alpha) = \inf_{[\Gamma]=\alpha} A(\Gamma)$. This is a typical minimax approach.

Theorem 7 *$W(\alpha)$ is continuous, and $\frac{dW}{dt}(\alpha) \leq -2\pi - \frac{1}{2} R_{\min}(t) W(\alpha)$.*

It's exactly the same as before except I have 2π, not 4π. You have to be careful about the boundary, you flow it by the curve-shortening flow, and that keeps you from needing a boundary correction term.

I don't actually work with forward difference quotients, I work on a fixed small interval.

Suppose, fix t_0, t_1. I want to say that for every loop, the area of the minimal disk of the loop at time t_1 is at most $V_{A(c(t_0))}(t_1)$. If I can show this for each loop, I can substitute in the solution for the biggest value. So $A(\Gamma)$ at time one is below the solution for $A(\Gamma)$ at time zero. So taking limits that's true for any time in this interval.

Curve-shortening in a three-manifold can develop singularities. You cross with a really tiny circle and study what they call ramps. You make it a graph over that tight circle. On ramps it's very good. All the estimates don't care about the tightness of the circle, they only care about the curvature, so you can take a limit back.

This is nontrivial. There's no boundary issues the other way, but they get index one critical points in the energy functional. We pay with something in dimension one, which is easier to understand. Notice you get the same estimate on each piece.

This argument does not say anything about acyclic components. This might bubble off three-spheres over and over again. There's no differential equation coming up from below telling you how this will disappear. You can say how long it takes now that you have the three-sphere, but nothing tells you you can't keep doing this, keep bubbling things off.

And we've proved the Poincaré conjecture.

I would like to thank the organizers, I think this was a great conference, I think we should thank Peter and Tomasz in absentia.