The Poincaré Conjecture

Colin Rourke

The Poincaré Conjecture states that a homotopy 3–sphere is a genuine 3–sphere. It arose from work of Poincaré in 1904 and has attracted a huge amount of attention in the meantime. It can be generalised to all dimensions and was proved for $n = 2$ in the 1920s, for $n \geq 5$ in the 1960s and for $n = 4$ in the 1980s. The Perelman program for proving the original case of $n = 3$, following earlier work of Hamilton, has been the subject of intense scrutiny over the last four years with several authors working hard on details. The prognosis is good, but it is probably still too soon to state categorically that the conjecture has been proved.

In this essay I shall discuss all cases of the generalised conjecture and outline the proofs for dimensions 4 and higher. Then I shall turn to Thurston's Geometrisation Conjecture (which implies the Poincaré Conjecture) and the Perelman program for proving it.

57M40; 57N10, 57R60, 57Q30, 57R65

1 Introduction

Near the end of his *Cinquième Complément à L'Analysis Situs* published in 1904, Henri Poincaré asks the following question:

> *"Est-il possible que le groupe fondamental de V se réduise à la substitution identique, et que pourtant V ne soit pas simplement connexe?"*

[9, page 498, lines 14–15]. From the context V refers here to any closed 3–manifold and note that by *simplement connexe* Poincaré means

> *"simplement connexe au sens propre du mot, c'est-à-dire homéomorphe à l'hypersphere"*

[9, page 436, lines 2–3]. Thus Poincaré is asking: *Can a closed 3–manifold have trivial fundamental group and yet not be homeomorphic to the 3–sphere?*

Since Poincaré's time, "simply-connected" has become synonymous with "trivial fundamental group" and the negative answer to Poincaré's question—*every closed simply-connected 3–manifold is homeomorphic to the 3–sphere*—has come to be known as the **Poincaré Conjecture**, though it is important to note that Poincaré himself made no guess about the answer to his question.

Poincaré's work on the subject is interesting. He originally announced that any homology sphere (ie a 3–manifold whose first homology group is zero) is necessarily the 3–sphere, but then discovered that this is false. He found the famous Poincaré sphere, now usually known as dodecahedral space, which is a homology sphere with a fundamental group of order 120.

In the intervening 103 years, the Poincaré Conjecture has elicited a huge amount of work. It is now seen as a central problem in the understanding of 3–manifolds. Many proofs have been announced by famous (and otherwise highly reliable) mathematicians, only to be withdrawn when difficulties in the proofs became apparent. In 2003, Grisha Perelman placed a detailed sketch of a proof on a public preprint website [8]. This proof follows naturally from earlier work of Richard Hamilton, indeed can be regarded as the completion of Hamilton's program for proving it. So far this proof appears to be an exception: although many difficulties have been found, it has proved possible to overcome them all. There is not yet a broad consensus that the proof is complete and correct, but the prognosis is very good and it seems likely that the conjecture has finally been proved.

In this essay, I shall first set the conjecture in its context by discussing the problem of classifying manifolds in general and then discuss the generalisations of the conjecture to higher dimensions, which have all been proved. Finally I shall come back to dimension three and discuss Thurston's Geometrisation Conjecture, which sets the Poincaré Conjecture in a wider context, and the Perelman program for proving Thurston's conjecture.

2 The classification of manifolds

Topology is the study of spaces under continuous deformation (the technical term is homeomorphism). It arose from Poincaré's work at the end of the nineteenth century. It has been one of the major strands of mathematics throughout the twentieth century and the start of the twenty-first.

The most important examples of topological spaces are manifolds. A manifold is a space which is locally like ordinary space of some particular dimension. Thus a 1–manifold is locally like a line, a 2–manifold is locally like a plane, a 3–manifold like ordinary 3–dimensional space and so on. In order to avoid pathological examples, it is usual to insist that a manifold has some additional structure, which tames it; the principal extra structures which are assumed are (a) differential (Diff) structure—the ability to do analysis on the manifold—and (b) piecewise-linear (PL) structure, which is roughly equivalent to assuming that the manifold can be made of straight pieces. In low dimensions (less than seven) these two seemingly opposed structures are in fact equivalent, and we shall freely assume either of them as necessary.

Let us explore manifolds by starting in low dimensions and working upwards. In dimension 1, there is very little to discuss. If we assume that our manifold has a PL structure, then it must consist of a number of intervals laid end to end. If we also assume that it is connected (all in one piece) then it either comprises an infinite collection of intervals forming a line or it comprises a finite collection forming a circuit. Thus a connected 1–manifold is homeomorphic either to a line or to a circle. Of these only the circle is compact (a technical notion equivalent to being made of a finite number of straight pieces). So there is just one compact, connected 1–manifold up to homeomorphism.

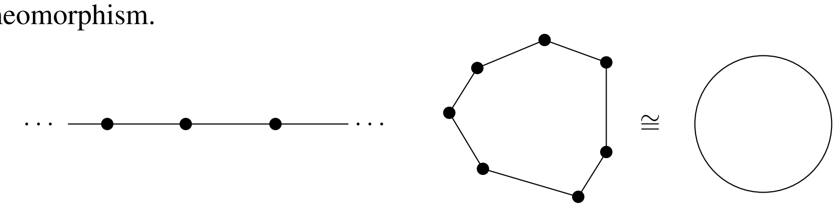

Moving on to dimension 2, one of the early successes of topology was the complete classification of compact connected 2–manifolds up to homeomorphism. It was proved in the 1920s that there is a simple list of such manifolds, which we shall describe below. It fact the classification of PL surfaces was given by Dehn and Heegaard in 1907. It was not until 1925 that a proof was given (by Rado) that any 2–manifold can be given a PL structure, thus completing the classification.

The simplest 2–manifold is the 2–sphere (the surface of a ball).

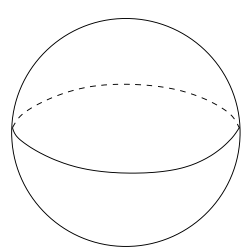

The next simplest and easily visualised is the torus (the surface of an american doughnut).

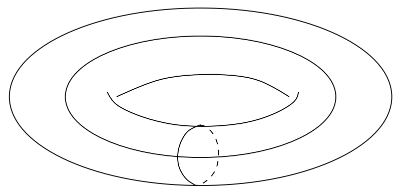

An infinite family of surfaces can now be described. Join two tori together by a tube to form a double torus. (This joining operation is called *connected sum*.)

(2–1)

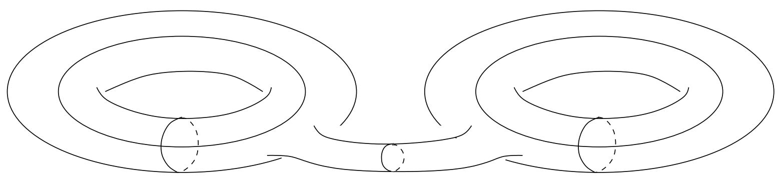

Now repeat to form a sequence of multiple tori (or pretzels):

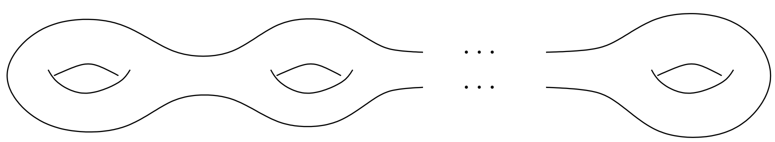

There is another infinite family of 2–manifolds, which is not quite so easy to visualise. Start with the projective plane (or cross-cap). The simplest description of this, is as a disc with opposite points on the boundary glued together:

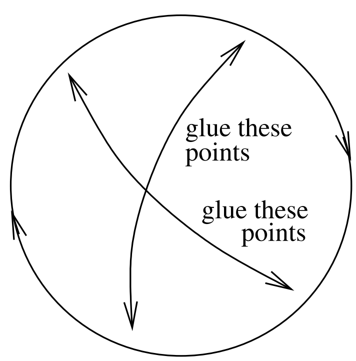

The resulting surface cannot be embedded in 3–space, but it can be immersed, that is locally embedded. The nicest immersion is the famous Boy's surface (here in the form of a sculpture installed in the Mathematishes ForshungsInstitut Oberwolfach):

Another nice description is to take a Möbius band. This is a surface with a boundary which is very easy to visualise:

(2–2)

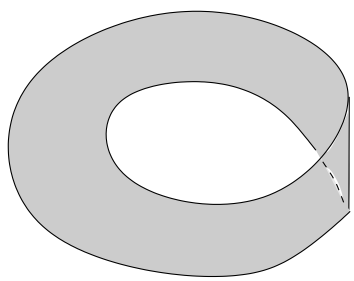

Now take the boundary of the band and cap it with a disc to form a surface without boundary which is in fact the projective plane again.

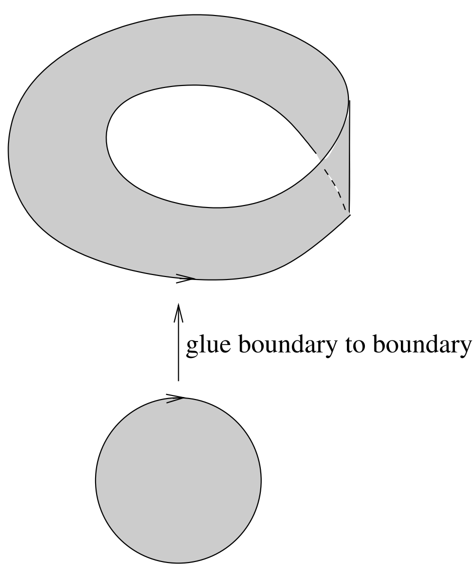

The second infinite family of surfaces is obtained by forming a repeated connected sum of projective planes.

(2–3)

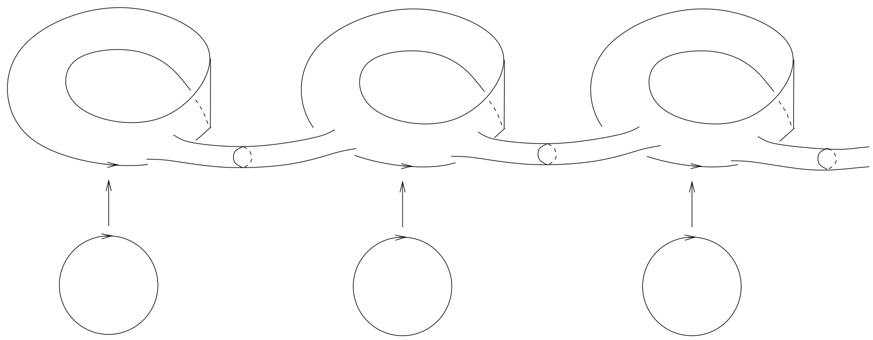

The classification theorem for 2–manifolds can now be stated. Any compact connected surface is homeomorphic to one of these two infinite lists. No two surfaces in the lists are homeomorphic.

This is the perfect classification theorem. We have a complete list and a simple description of any surface in the list. It is a theorem of this type that we would like to have for all dimensions.

Moving up to dimension 3, there may be a similar (but rather more complicated) complete list of manifolds. There is a consensus of what should go on the list. But it has not been proved. Indeed one of the major ingredients of the proof would be a proof of the Geometrisation Conjecture, which implies the Poincaré Conjecture. There has been a lot of recent progress with classification assuming the Geometrisation Conjecture which now remains the major step in the program. We shall discuss this conjecture and the Perelman program in detail in the last two sections of this essay.

Moving up further to dimension 4, it has been proved that there is *no* complete list of manifolds similar to the list in dimension 2 or the conjectured list in dimension 3. This is for a simple, but highly technical reason. It is possible to construct a well-defined 4–manifold with any given finite presentation for its fundamental group. Now it is known that there is no algorithm which will decide, given two group presentations, whether the groups are isomorphic. A homeomorphism between two manifolds induces an isomorphism of fundamental groups. It follows that it is not possible to decide whether two 4–manifolds are homeomorphic.

Similarly, there is no hope of classifying manifolds of dimension bigger than 4. However the Poincaré Conjecture does not deal with a general manifold and so, as we shall shortly see, the failure of classification does not imply the failure of proof of the Poincaré Conjecture.

For basic topological background, including the fundamental group (outlined in the next section) and for further reading on the classification of 2–manifolds, I recommend Armstrong [1].

3 The conjecture and its generalisations

It is time to examine in detail what the conjecture actually says in all dimensions. We gave the statement for $n = 3$ in Section 1. Let us start by examining this statement in greater detail. It turns on the fundamental group. This is an invariant constructed by Poincaré. What we do is to consider loops in the manifold. A loop is just a path that starts and ends at a chosen basepoint. We add loops by concatenating them, ie by making a new path which is the first path followed by the second. This addition makes the set of loops into a group. This is a bit of a simplification. We need to allow loops to vary continuously (called homotopy). Then the homotopy classes of loops under loop addition defines the fundamental group. Having trivial fundamental group is the same as having the property that every loop is shrinkable (homotopic to the constant loop).

Next we need to define spheres in each dimension. A 1–sphere is just a circle. A 2–sphere was described above. It is the surface of a ball. In each case we could have defined the sphere to be the points in ordinary space (Euclidean space) at unit distance from the origin. Thus the circle comprises all points in the plane (2–space) at unit distance from the origin. The 2–sphere is all points in ordinary 3–space at unit distance from the origin. We can now generalise. The n–sphere comprises all points in $(n + 1)$–space at unit distance from the origin. We use the notation S^n for the n–sphere. We have defined the 3–sphere (S^3) using 4–space. However it is a mistake to think of S^3 as 4–dimensional. It is in fact just ordinary 3–space with an extra point thrown in. The way to think of this is to think about maps of the earth. There is a map called "stereographic projection" which maps the surface of the earth onto a plane, with the north pole omitted (it goes to infinity). Reversing this map, we can think of the surface of the earth (ie a 2–sphere) as the plane with just one point added at infinity. In a similar way we can think of S^n, for general n, as n–space with just one point added. So S^3 really is a 3–dimensional object.

The Poincaré Conjecture is that a (compact connected) 3–manifold whose fundamental group is trivial is homeomorphic to S^3.

For 2–manifolds (surfaces) the conjecture makes sense in exactly the same terms: every (compact connected) 2–manifold whose fundamental group is trivial is homeomorphic to S^2. Indeed looking at the list of 2–manifolds given earlier, it is easy to see that all have non-trivial fundamental groups apart from S^2. So this is in fact a theorem not a conjecture. However if we move up to dimension 4 then it is easy to construct compact connected 4–manifolds with trivial fundamental group which are not the 4–sphere. One simple example is "complex projective space", which is the complex analogue of the projective plane described earlier. Since this manifold plays no further role in this essay, I shall say no more about it.

Moving up to higher dimensions, it is again easy to find compact connected manifolds with trivial fundamental group which are not the sphere of the appropriate dimension. Therefore to generalise the conjecture to higher dimensions sensibly we need to change it to an equivalent conjecture. The idea of "homotopy equivalence" is the key one.

The concept of homeomorphism was touched on briefly before. A homeomorphism between two spaces is a continuous bijection between them with continuous inverse. The idea is that it may change sizes—it may deform the space—but it may not tear it apart. This is what continuity means.

The idea of homotopy equivalence is to allow general maps rather than bijections and ask that the compositions be merely homotopic to bijections rather than actually being bijections. The homotopy allows great flexibility and makes the relationship much more general.

For a simple example, consider an annulus. This can be squeezed down to a circle:

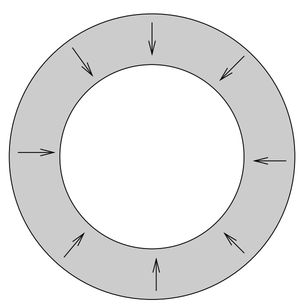

This squeezing is not a homeomorphism but it is a homotopy equivalence. In a similar way a Möbius band (figure 2–2) can be deformed to a circle by squeezing onto the centre line. Thus all three (Möbius band, annulus and circle) are homotopy equivalent. Homotopy equivalence captures the global algebraic invariants of a space, whilst losing the local geometry.

Let us call an n–manifold a "homotopy sphere" if it is homotopy equivalent to S^n. Now it can be proved that a 3–manifold is a homotopy sphere if and only if it is simply-connected (has trivial fundamental group). Thus the Poincaré Conjecture can be reformulated:

Every homotopy 3–sphere is homeomorphic to S^3.

This statement can easily be generalised to any dimension. Further there are now no easy counterexamples:

Generalised Poincaré Conjecture *Every homotopy* n*–sphere is homeomorphic to* S^n.

As remarked above this is a true theorem for $n = 2$. In the next two sections we discuss the cases $n \geq 5$, which were proved in the 1960s, and the case $n = 4$, which was proved in the 1980s. After that we return to dimension 3 and the Perelman program for the proof of the original conjecture.

4 High dimensions: $n \geq 5$

It may seem paradoxical, but the generalised conjecture is far easier to prove in high dimensions. It is not hard to explain the proof in general terms and I shall attempt to do this. There were two independent lines of proof,

(1) using the idea of *engulfing* (independently due to Zeeman and Stallings, using somewhat different formulation), and

(2) using *handle theory* (Smale).

The engulfing theorem in its most powerful form—suitable for proving all the cases $n \geq 5$—was first worked out by Zeeman. The Stallings and Smale approaches were initially limited to $n \geq 7$. However all three methods can be used to prove all cases $n \geq 5$ and it is usual to credit all three of these mathematicians (Smale, Stallings and Zeeman) with these results.

The Smale approach is far more powerful. Using handle theory it is possible to recover engulfing. Handle theory proves stronger versions of the generalised conjecture: for all three approaches it is necessary to assume that the homotopy sphere has a PL or Diff structure. Using handle theory it can be proved that the homotopy sphere is homeomorphic to the genuine sphere preserving this extra structure (to be precise, this is true for PL structures in all dimensions ≥ 5 and for Diff for $n = 5$ or 6: for $n \geq 7$ there are smooth homotopy spheres constructed by Kervaire and Milnor which are homeomorphic but not diffeomorphic to genuine spheres). Using engulfing, all that can be deduced is a homeomorphism without the extra structure.

However, for historical interest, and because it is easy to do so, I shall outline the engulfing approach first, following a suggestion of Zeeman in the late 1960s which gives a great simplification.

The idea of the proof by engulfing is to cut the homotopy sphere M into complementary pieces which contain the parts of the PL structure up to about half the dimension (the technical name for these is "skeleta" and one of these uses the "dual" skeleton) and then prove (using the fact that M is a homotopy sphere) that each of these pieces is contained in a ball, ie *engulfed* by the ball, where a *ball* means a homeomorphic copy of half a sphere. Then by a simple pushme–pullyou argument we can assume that these balls cover M. Push a neighbourhood of X which is inside the ball engulfing X out until it is close to X'. Then pulling it back pulls the ball engulfing X' to cover everything not covered by the one engulfing X. It then follows by using a result known as the Schoenflies Theorem that the manifold is in fact a sphere. So the guts of the proof is to prove that a skeleton X of M of about half the dimension is contained in a ball.

Since M is a homotopy sphere we can find a deformation of X to a point in M. This defines a map of the cone $C(X)$ on X into M:

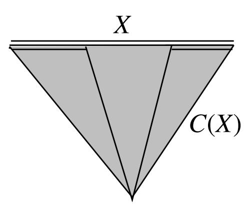

The cone is a concrete realisation of the deformation of X to a point. We can deform down the cone lines in a systematic way known as collapsing. If the map of $C(X)$ does not meet itself (we say it is non-singular) then the collapsing process can be reversed to pull a little ball at the cone point back over the whole of $C(X)$ and we have engulfed X. However, we cannot expect the map to be non-singular. At this point we use *general position* which is a simple technique to reduce the dimension of the singular set as far as possible. For example suppose that we have a map of the circle into the plane, then we expect the self-intersections to be isolated points.

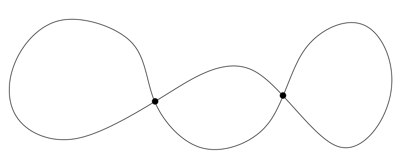

Similarly two planes in 3–space meet each other in a line in general position. In general position a map of a p–dimensional object into an m–manifold has singular set (where it meets itself) of dimension $2p - q$. Now assume that M has dimension m and that X has dimension $\leq m - 4$ (we say X has *codimension* ≥ 4) then $C(X)$ has codimension ≥ 3 and S the singular set has dimension $\leq 2m - 6 - m = m - 6$. Now the collapse of $C(X)$ to the conepoint induces a collapse of the image of $C(X)$ in M to the *shadow* X' of S in $C(X)$:

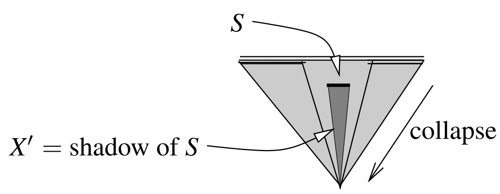

Now X' has dimension $\leq m - 5$ and if we can engulf X' then by reversing the collapse we can engulf X. So we have replaced the problem of engulfing X by engulfing X' which has smaller dimension. We repeat this argument reducing dimension at each stage until we get to dimension 0 when engulfing is obvious.

If M has dimension ≥ 7 then this argument proves the conjecture. We choose for X the 3–skeleton of M and let X' be the dual $(m - 4)$–skeleton. These suffice for the pushme–pullyou argument, and since each can be engulfed in a ball, we are done. To make the argument cover $n \geq 5$ we need to work with subsets of codimension 3 rather than 4. In this case the argument stalls at the first stage beacuse X' could be the same dimension as X instead of smaller. This is where the Zeeman trick comes in. The top dimensional singularities come in pairs σ_1, σ_2 which get identified in the image, as indicated at the left in the diagram below. What we do is to find subcones C_1, C_2 containing σ_1, σ_2, then remove the inside of C_2. At this point C_1 collapses. After the collapse we put back C_2 and collapse it in turn. The result is Y. If we can engulf Y then running the process backwards, we can engulf X. But the singularities in Y are one dimension lower and the argument used in codimension ≥ 4 applies to engulf it. The whole argument is encapsulated in the following diagram:

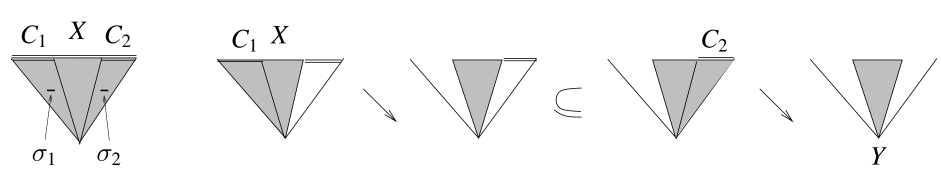

Now let us look briefly at Smale's handle theory argument.

The idea of a handle decomposition is to describe a manifold as a union of standard objects. This is best understood by example:

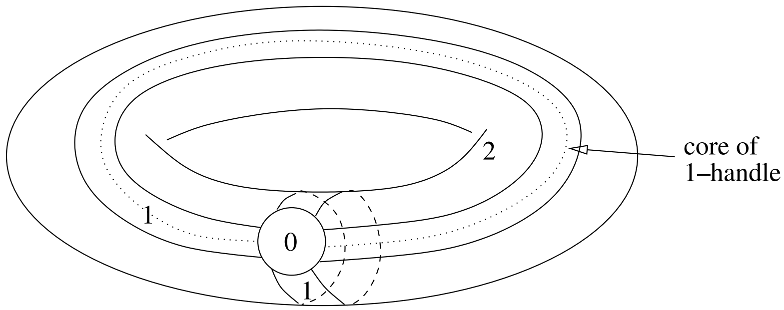

The figure shows a torus made by starting from a small disc (0-handle) and attaching two bands (1–handles) and finally filling in the resulting hole with a big disc (2–handle). Each handle is a disc, but the way of thinking of them varies according to the index. For example we think of the 1–handles as neighbourhoods of their 1–dimensional "cores" which are the central intervals. It is a standard result in differential (or PL) topology that any manifold can be given such a decomposition.

The Smale proof is based on the idea of an "h–cobordism" which is a manifold W with two boundaries (ends) each homotopy equivalent to the whole manifold.

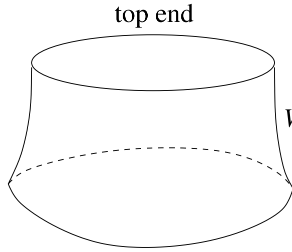

The h–cobordism theorem states that if dimension is at least six then any h–cobordism is trivial (ie a product with an interval) and in particular the two ends are homeomorphic. To prove the Poincaré Conjecture from the h–cobordism theorem, remove two balls from the homotopy sphere. The result is an h–cobordism.

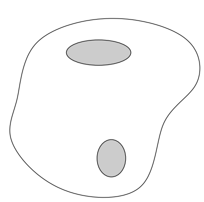

If it is trivial then it can be seen that the homotopy sphere is a genuine sphere. This argument works for dimension ≥ 6. For dimension 5, a preliminary argument is used to construct an h–cobordism from the homotopy sphere to a standard sphere.

The method of proof of the h–cobordism theorem is to read the homology groups of the manifold from the decomposition. Since the inclusions of the ends are homotopy equivalences the relative homology groups all vanish. Reading this back to the handle decomposition, we can prescribe a number of simple moves which will trivialise the decomposition (and prove the h–cobordism theorem). The key move is "handle cancellation". If two handles of adjacent index are arranged like this, with the attaching sphere of the second handle meeting the belt sphere of the first in a single point then they cancel.

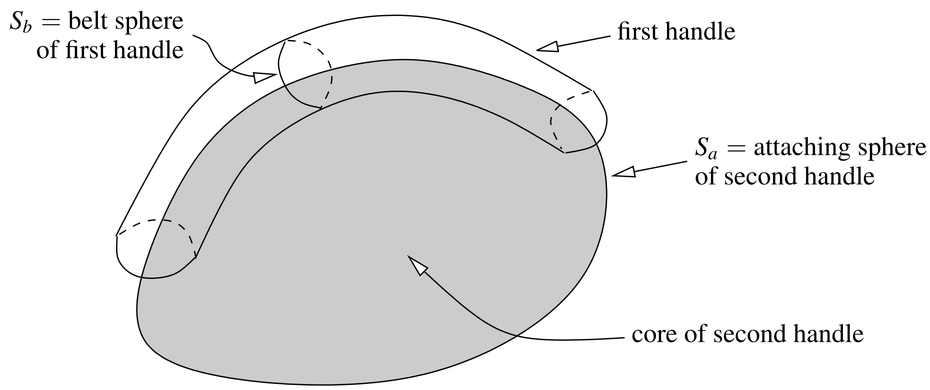

The handles are said to be "complementary". The cancellation is given by "pushing" the first handle down over the second. However reading the homology from the decomposition just tells us that the handles are *algebraically* complementary, which means that the points of intersection of the belt and attaching spheres (S_a and S_b in the figure) can be paired off apart from one essential intersection. The crux of this part of the argument is to cancel the extra pairs of intersections. This is done by using a result known as the Whitney trick. We join the points by arcs α, β in the two spheres, then span the circuit formed by the two arcs by a disc (the Whitney disc) and then push S_b across the disc to eliminate the two points of intersection.

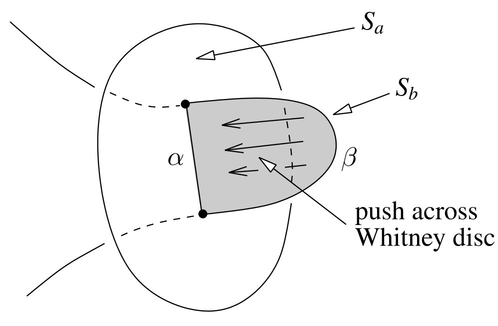

Once all the unnecessary points of intersection have been removed, the handles become complementary and cancel.

The Whitney trick works nicely if the dimension of the ambient manifold is at least 5, when the required disc exists by general position. In the next section I shall sketch how the proof was modified for the case of 4–manifolds.

For further reading on the proofs of the higher dimensional Poincaré Conjecture, I recommend Buoncristiano and Rourke [2] for the engulfing argument and Milnor [6] (Diff case) or Rourke and Sanderson [10] (PL case) for the Smale proof using handle theory.

5 Dimension 4

The case $n = 4$ of the Generalised Poincaré Conjecture took a further 20 years to be proved. The eventual method of proof was similar to the Smale proof in dimension ≥ 5, using handle theory methods, but the extension was highly non-trivial. The problem is the Whitney trick. If we try to push the argument outlined above down to dimension 4, then the Whitney disc D is not obviously embedded. It may have double points of its own. So the argument simply trades the pair of points of intersection of S_a and S_b, which we need to eliminate, for one or more double points in the Whitney disc. Casson's brilliant idea was to repeat the argument infinitely many times. Consider a double point P in D and draw a curve in D from one of the sheets at P round to the other forming a loop α_1 which in turn bounds a disc D_1 (figure below). Now D_1 may have further double points. Find such a disc D_2, D_2' etc, for each double point in D_1 and then do it again for each double point in each disc D_2, D_2' and so on, and keep repeating this for each new disc that we use. The resulting infinite construction is a tower of discs with double points indicated in the figure.

The neighbourhood of the Casson tower in the 4–manifold we are considering is called a *Casson handle*. Casson's hope was that the problem of finding an embedded version of the original disc D might somehow be pushed to infinity by this construction. It is easy to see that the fundamental group obstruction is killed by the construction but this is a far cry from proving that there is an embedded disc. Casson's hope was brilliantly vindicated by Freedman (about 10 years later) who proved, in an argument too intricate to give any real flavour of here, that a Casson handle is a genuine topological handle. Then the Whitney disc can be found and the h–cobordism theorem, and hence the Poincaré Conjecture, proved. Freedman's argument uses a style of topology which goes

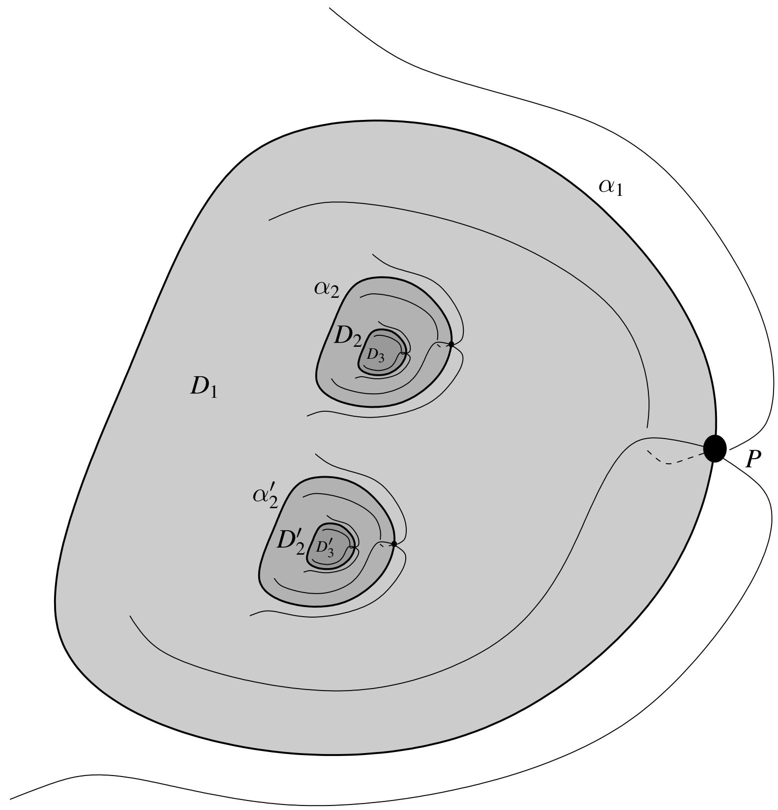

back to R L Moore and R H Bing. Freedman referred to the proof as "my take-home exam in Bing topology".

Before leaving dimension 4 and moving to dimension 3, it is worth commenting that Freedman's methods are strictly topological. They do not work with either a Diff or a PL structure. Indeed the Diff and PL versions of the 4–dimensional Poincaré Conjecture are still unproved. Unlike dimensions 2 and 3, where there is an equivalence between topological manifolds and Diff and PL manifolds, in dimension 4 the theories diverge drastically. A Casson handle may not contain a *smooth* Whitney disc. Furthermore there is only one topological structure on ordinary 4–dimensional space whilst there are uncountably many Diff (or equivalently PL) structures. These bizarre facts are a consequence of work started by Donaldson and continued by many other authors, but it is a subject too far removed from the Poincaré Conjecture to outline further here.

For further reading I recommend Freedman and Quinn [4] for the topological case outlined here and Donaldson and Kronheimer [3] for information on the differential topology of 4–manifolds.

6 Thurston's program

We now return to dimension 3. The excursion we have made into higher dimensions was worth the time if only because it gives a good flavour of the kind of argument that mathematicians found impossible to push through in dimension 3. The methods that appear likely to have solved the 3–dimensional problem are, as we shall see, quite different.

However, one idea that works well in other dimensions also works well in dimension 3, namely the idea of connected sum. An early result (1929) of Kneser proves that any 3–manifold is the connected sum of *prime* 3–manifolds, ie 3–manifolds which are not themselves connected sums (a corresponding uniqueness result was later proved by Milnor). The picture is similar to the one for surfaces given earlier (figures 2–1 and 2–3) but now the connecting tubes have cross-section a 2–sphere instead of a circle (1–sphere).

(6–1)

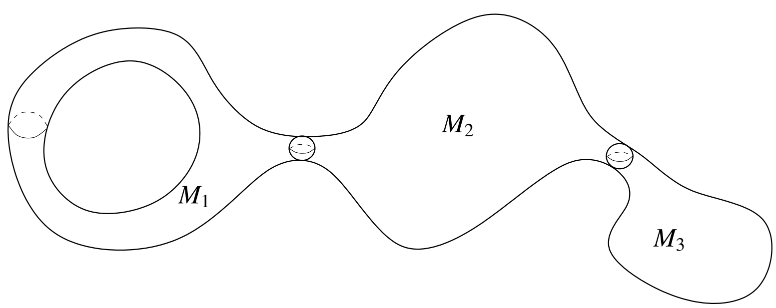

For 3–manifolds there is another finer (and again unique) decomposition given by cutting along tori and Kein bottles. This is the JSJ (Jaco–Shalen–Johannson) decomposition.

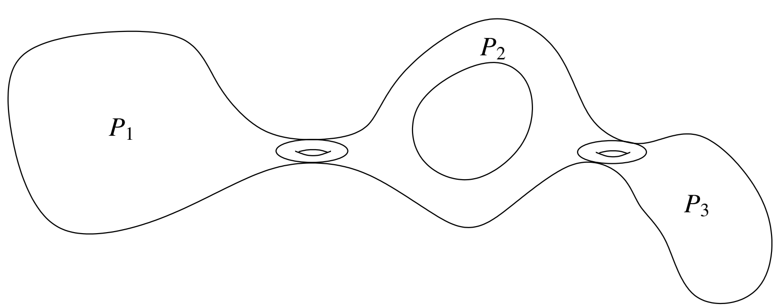

Let us call the pieces into which the 3–manifold is cut by both decompositions the "elementary pieces". Thurston's Geometrisation Conjecture can now be stated:

Geometrisation Conjecture (Thurston) *Each elementary piece admits a unique geometric structure.*

To understand what this conjecture says, and what it implies for the Poincaré Conjecture, we need to talk a little about what a geometric structure is. It is best to go back to dimension 2. The ordinary 2–sphere has a natural geometry. We can measure distance along great circles, measure angles, prove theorems about triangles and so on.

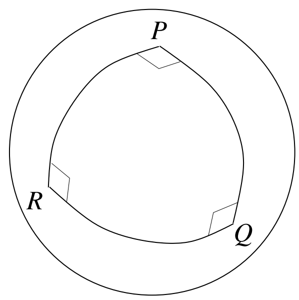

Topologically the surface of any bumpy or squashed ball is a 2–sphere. Only the perfectly round ones are geometric. Triangles on a (round) 2–sphere have angle sum bigger than π; see for example triangle PQR in the figure, which has angle sum $\frac{3\pi}{2}$; we say that the sphere has *positive curvature*. Moving up to the torus, the usual embedding in 3–space is not geometric. The torus is positively curved at the outside and negatively curved inside.

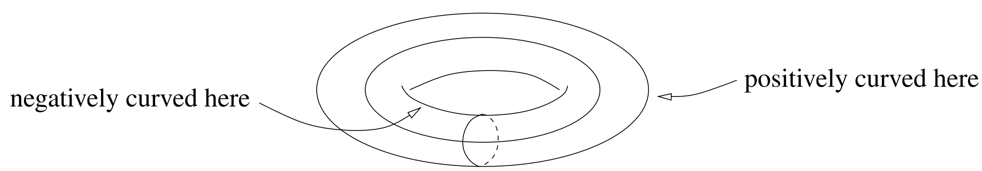

To be geometric, each pair of points must be equivalent and hence the curvature must be uniform. But there is a way to think of the torus as obtained by gluing opposite sides of a square together.

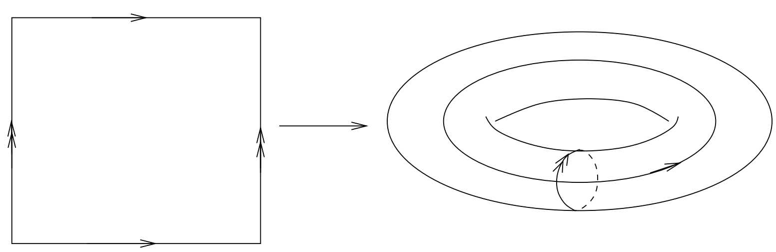

The square is part of the ordinary plane, which has a familiar geometry (Euclidean geometry). As we cross the glued sides of the square, the geometry carries over unchanged and thus we have a natural geometric structure on the torus. Triangles in this geometry have angle sum exactly π and we say that this is a *flat geometry* or that the torus has *zero curvature*. Moving on to the double torus, we can again regard it as a polygon (in fact an octagon) with sides glued, but to get the angle right at the vertex, we need to use a non-Euclidean octagon with all angles $\frac{\pi}{4}$. Then the non-Euclidean geometry of the octagon carries over to give a geometric structure to the double torus.

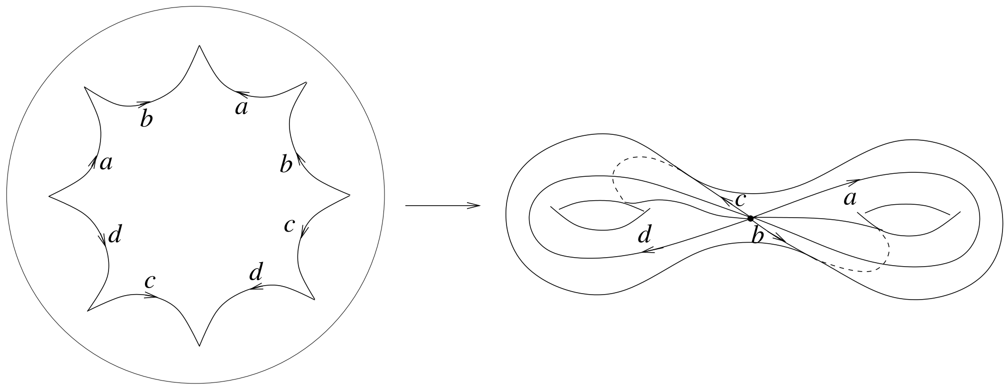

In the figure we have shown the octagon as part of the Poincaré disc model for non-Euclidean geometry. Triangles in this geometry have angle sum less than π and we say it has *negative curvature*. A similar (rather more complicated) construction works for multiple tori, which all admit geometric structures of negative curvature.

There is an analogous way to put geometric structures on connected sums of projective planes. The projective plane itself can be thought of as the 2–sphere with opposite points identified and so has a natural geometry of positive curvature. The connected sum of two projective planes is a Klein bottle, which has a flat geometry by a similar construction to the torus, and connected sums of three or more projective planes have natural geometric structures of negative curvature.

The best way to think of the Geometrisation Conjecture is as the 3–dimensional analogue of this geometric description of surfaces. There are three geometries which are analogous to the three used in dimension 2, namely:

- Spherical: the geometry of S^3 (positive curvature)
- Flat: the geometry of ordinary 3–space (zero curvature)
- Hyperbolic: the geometry of non-Euclidean or hyperbolic 3–space (negative curvature)

There are five other geometries that are needed to geometrise 3–manifolds but for understanding the Perelman program as it applies to the Poincaré Conjecture, we can safely push them to the background. All we need to know is that any elementary piece which has a geometry other than spherical must have an infinite fundamental group. This follows from some elementary covering space theory. The universal cover (a well-defined simply-connected space which can always be constructed for manifolds and other locally well-behaved spaces) is a copy of the archetype of the geometry (eg S^3 for spherical, $\mathbb{R}^3$ for flat etc) and all except S^3 are infinite, which implies that the fundamental group of the manifold is infinite.

So now suppose that we have a prime homotopy 3–sphere. By the Geometrisation Conjecture it has a geometric structure and by the remark just made it must be spherical. Then the universal cover is a copy of S^3 as remarked above, but being simply-connected it is already its universal cover, in other words it is itself a copy of S^3. So the Geometrisation Conjecture implies the Poincaré Conjecture.

Before moving on to the Perelman program, it is worth remarking that progress on the classification of geometric 3–manifolds (and therefore of all 3–manifolds if the Geometrisation Conjecture is proved) is very advanced. There is a complete classification of geometric 3–manifolds for all geometries except hyperbolic. The classification of hyperbolic 3–manifolds is currently an intense research area.

For further reading I recommend Thurston's notes [11].

7 Perelman's program

The Perelman program is based on work of Hamilton, who introduced the idea of the *Ricci flow* on a differentiable manifold. We assume that our manifold has a *metric*, ie a precise way of measuring distances. The Ricci flows acts on the metric. The idea is that it should smooth out the irregularities in the metric and result in a perfectly regular metric, in other words a geometric structure. Before wading into the technicalities of this program it is worth remarking that this technique is quite different from the techniques that proved the higher dimensional Poincaré Conjecture and which we outlined in sections 4 and 5. In the engulfing proofs we considered nice subsets of the manifold and operated on them, and in the handle theory arguments we cut the manifold up into nice pieces and then moved them around. The proofs work by passing into the detailed makeup of the manifold. By contrast the Hamilton–Perelman program is holistic. It takes the entire manifold and operates on it all at once.

The Ricci flow is defined by a partial differential equation on the space of metrics on a given manifold M. Here is the equation which defines the flow:

$$\frac{\partial g}{\partial t} = -2\mathrm{Ric}_g \tag{7–1}$$

In the equation g denotes the metric and Ric_g denotes the *Ricci curvature* of the metric. The equation has to be understood as an evolution equation. We consider all possible metrics on our manifold and we let the metric evolve so that the rate of change at any point is $-2\mathrm{Ric}_g$.

To understand this term, we need to talk about *curvature*. We mentioned it in the last section in connection with geometric surfaces. To decide how curved a surface is, look at triangles drawn on it and compute the angle sum. One way to detect a triangle with angle sum different from π is to transport a vector around it. Below we have illustrated the result of transporting a vector around a triangle on a sphere. The resulting vector is not parallel to the starting vector. This corresponds to the fact that the angle sum of the triangle is bigger than π.

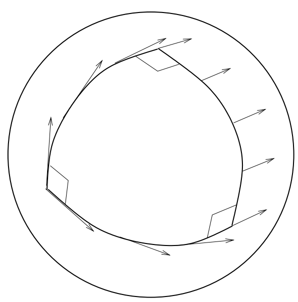

It is not neceesary to use a triangle to detect curvature. Transporting a vector around any closed curve on the sphere results in a non-parallel vector. The same idea can be used in any manifold with a metric. By transporting vectors around small curves we can define curvature. If we choose our curve to lie in a plane, we get the notion of the curvature of that plane. But notice that this is a vector—the discrepancy after transport—not a number. To define the Ricci curvature, choose a fixed vector and average curvature over all planes containing that vector. Unlike the 2–dimensional case, where there is just one local curvature, the Ricci curvature in dimension three is a "tensor" with 9 components (each curvature has three components and there are three independent directions to use). But there is some symmetry and there are in fact just 6 independent components. The diagonal components of the Ricci curvature are the sectional curvatures which we can think of as directly analogous to the curvature of a surface; these are the curvatures of three mutually perpendicular hyperplanes measured in a perpendicular direction. If we add these three we get the *scalar curvature* which gives a rough idea of how curved the manifold is. It was Hamilton's insight, vindicated by Perelman, that the Ricci curvature is the correct notion of curvature to use to evolve a general metric.

Perhaps the best way to try to understand the Ricci flow (equation 7–1) is by analogy with other, more intuitive, flows. The *heat equation* describes the diffusion of heat in a solid body. we have a very clear idea of how heat flows. It evens itself out. Highs get lower, lows get higher and the (asymptotic) end result is a uniform distribution with each point having the same temperature. The Ricci flow equation describes a similar flow on the metric of a manifold: it smoothes out local variations. But there are obvious differences. The metric does not necessarily become uniform. If we start in a region where the Ricci curvature is negative then the equation says that the metric grows, in other words points get pushed apart and the curvature moves upwards towards flatness. Conversely if we start where the curvature is positive then points get pulled together and the curvature becomes more intense.

A remarkable early result of Hamilton on the Ricci flow is that, if all points have positive Ricci curvature, then the flow does even the curvature out perfectly. The manifold shrinks towards a round 3–sphere (or a manifold which is covered by a round 3–sphere). This proves the Poincaré Conjecture for manifolds which admit metrics of positive Ricci curvature. However we cannot expect anything so nice to happen for a general 3–manifold. A general 3–manifold does not admit a perfectly uniform metric (ie a geometric structure). We need to cut the manifold into pieces first. So if the Ricci flow is going to prove the Geometrisation Conjecture then it must naturally lead to singularities whose resolution causes us to cut the manifold into pieces. Hamilton knew this but he could not analyse all the cases of what can happen when the flow goes singular, and his program ground to a halt.

It is easy to see that the flow must go singular in finite time for many starting conditions. This is because there is an evolution equation for minimum of the scalar curvature. It must change to become greater. If it ever becomes positive, then it must go to infinity in finite time, in other words the flow must become singular.

It is here that Perelman takes over the story. His first breakthrough was a "no local collapsing" result which completely overcame the problems that Hamilton ran into. Using this result and some very sharp and careful analysis he controls closely what happens when the Ricci flow becomes singular. What he proves is that the singularities which appear all have a tubular region which is a 2–sphere cross a line. What he then does is to cut across and glue in caps to fill in the 2–sphere. This is called "surgery".

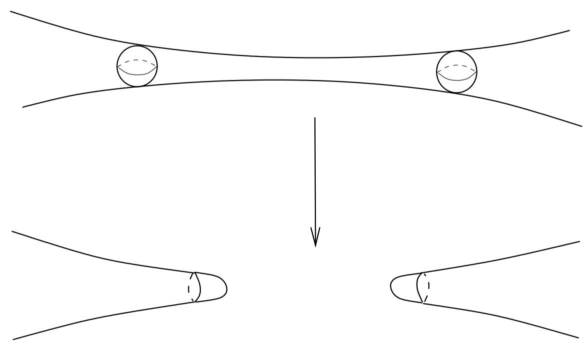

After surgery the Ricci flow can be restarted. The change is sufficiently simple that it can be described topologically. Indeed the topological picture is very similar to the connected sum decomposition pictured in figure 6–1. At this point, let us assume that we have started with a prime homotopy sphere. Then the surgeries can only remove trivial parts of the manifold and exactly one bit must be the manifold we started with. We now consider the sequence: flow ⟶ surgery ⟶ flow ⟶ surgery etc, which is called a *Ricci flow with surgery*. There are a number of careful estimates which can be made for the Ricci flow after surgery in terms of the flow before and we can analyse the whole sequence, just as if we were not changing the manifold. The crux of the proof now is to prove two key results:

(1) Surgery times do not accumulate: we are not forced to perform infinitely many surgeries in a finite time interval.

(2) The whole flow becomes extinct after a finite time.

It then follows that, after a finite number of surgeries which discard trivial parts of the original manifold, we arrive at a situation where the whole manifold shrinks to a point, which can be recognised as the end of the Hamilton proof for positive Ricci curvature. In other words we have arrived at a 3–sphere (or a manifold covered by a 3–sphere: but since we are assuming simple-connectivity, this latter case does not arise) and proved the Poincaré Conjecture.

The proof of (1) is one of the most delicate parts of the whole program. There are several parameters which are chosen each time we do a surgery. What Perelman proves is that these parameters can be chosen so that the amount of volume lost by the surgery (notice that some part of the original manifold gets cut out at surgery) is greater that a certain small amount. If we had infinitely many surgeries in a finite time, then an infinite volume would be lost. But it is easy to estimate the amount of volume gained by the flow and not enough is gained. So this proves (1).

To prove (2) Perelman uses an estimate for the change of area of minimal surfaces under Ricci flow with surgery, which is proved by an integral calculation. The area changes in a way which shows that area must decrease to zero in a finite time. If we could fill the whole manifold with such surfaces, then the whole manifold would shrink to zero, in other words the flow would become extinct.

It is not easy to fill a homotopy sphere with minimal surfaces. But we can fill it with curves and, for each curve, take a minimal disc which it bounds and this is what Perelman does. One of the properties equivalent to being a homotopy sphere is that there is a map of the genuine 3–sphere onto the manifold. There is a standard way to fill S^3 with curves, the Hopf fibration, pictured below. A half-plane is filled with circles as shown and these are rotated around the z–axis as indicated:

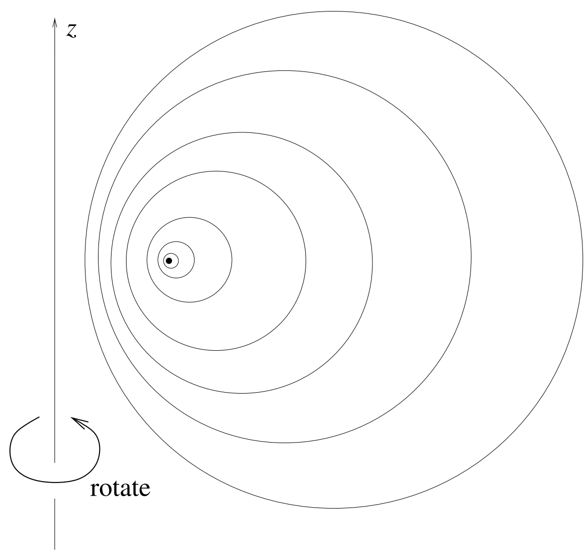

This fills 3–space with tori plus one central circle and the z–axis. Fill each torus with circles like this:

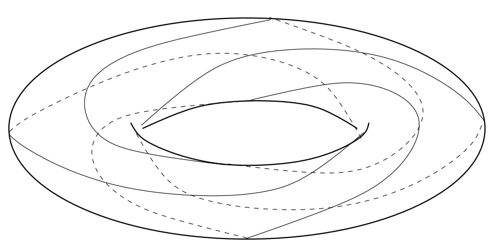

We have now filled the 3–sphere with the circles in the tori plus the central circle and the z–axis (which becomes a circle when the point at infinity is added).

If we now use the map of S^3 onto our homotopy sphere then the homotopy sphere is also filled by curves. For each such curve we find a minimal disc which it bounds. We let the Ricci flow act and we also let the curves flow using the *curve-shrinking* flow. This is a natural way to make curves in a manifold more uniform. Again we can think of it as analogous to heat flow: curves move uniformly towards curves of minimal curvature which are called *geodesics*. A similar integral calculation shows that the area of the minimal discs must then go to zero in finite time and again this calculation applies to Ricci flows with surgery.

However there is a serious technical problem to be solved. Like the Ricci flow, the curve-shrinking flow can run into singularities where the curvature of a curve goes infinite. Indeed the family of curves that we start with will very probably have such singularities in-built before the flow starts. Below we have sketched the generic singularity in a family of this type. It is called a *cusp* singularity.

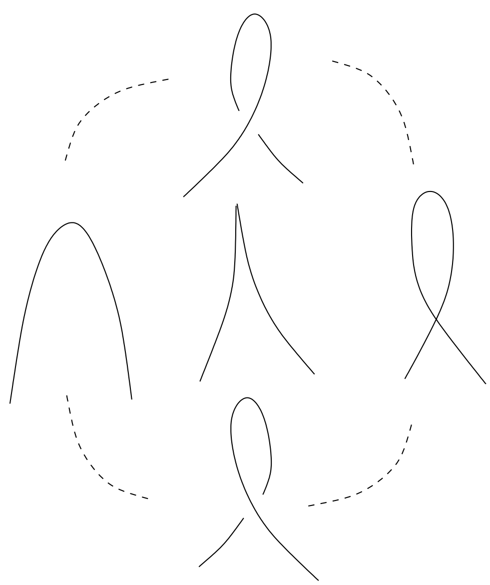

The diagram shows part of a 2–parameter family of curves in a 3–manifold. The central curve has a cusp; around it in a circle are curves which resolve the cusp in all possible ways.

To overcome this problem, Perelman uses a technique originally applied to plane curves (called "ramps") which, under the right conditions, allows the flow to pass through singularities. If we could use smooth general position at this point, then the singularities would be isolated (each curve would be singular for a finite number of times) and this technique would complete the proof. But general position does not apply to analytic flows of this type and instead Perelman uses the ramp technique, but not to the final limit. He stops near the limit and makes some very delicate estimates. These show that the time occupied by the singularities is small enough that the estimate, which shows that a minimal disc must shrink to zero in finite time, applies to all the discs in the family and they must all shrink to zero in finite time proving that the flow with surgery becomes extinct.

This completes the outline of the Perelman program for proof of the Poincaré Conjecture. For completeness I would like to finish by giving some idea of how the program proves the Geometrisation Conjecture. For this it is not possible to work only for a finite time. With a general 3–manifold, the Ricci flow continues indefinitely and indeed it is possible that there are an infinite number of surgeries (or rather it has not been proved that there are only a finite number). What Perelman does is to use similar estimates to those used to control the singularities to give a picture of the long-term behaviour of the flow. Eventually the flow settles down to give a decomposition of the manifold into a "thick" part and a "thin" part. To decide whether we are in a thick or thin part, we consider how far we can walk in all possible directions without coming back to our starting position. I have sketched this decomposition below.

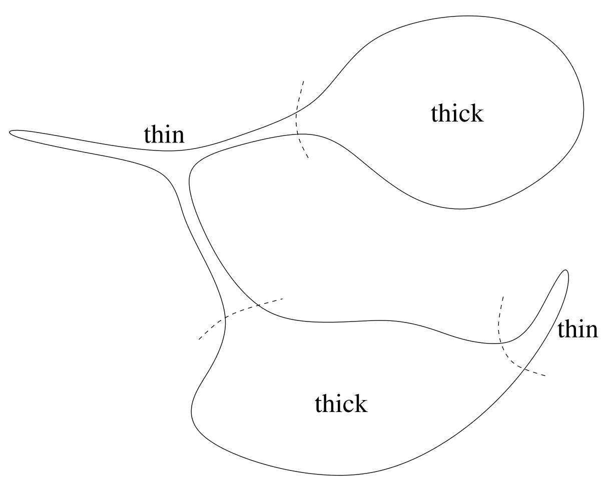

In the thin part, a short walk around one of the tubes gets us back to our starting position. The thick part has uniformly negative curvature, in other words a hyperbolic geometric structure. The thin part is what is know as a "graph" manifold. These manifolds can be described explicitly (they are connected sums of Seifert fibre spaces and surface bundles) and all have appropriate geometric structures; it is here that the geometries other than spherical and hyperbolic must be used. The final picture proves the Geometrisation Conjecture.

For further reading of an introductory nature on the Perelman program for the Poincaré Conjecture, I recommend the introduction to Morgan and Tian [7] and for the Geometrisation Conjecture, the introduction and outlines of Perelman's papers to be found in the notes of Kleiner and Lott [5]. For more detail on both topics, read the main texts of [7] and [5].

References

[1] M A Armstrong, *Basic topology*, Springer Undergrad. Text Ser. (1983)

[2] Sandro Buoncristiano and Colin Rourke, *Fragments of topology from the sixties*, Geom. Topol. Monogr. 6 (2003–7)

[3] Simon Donaldson and Peter Kronheimer, *The geometry of four-manifolds*, Oxford Math. Monogr. OUP (1990)

[4] Michael Freedman and Frank Quinn, *Topology of 4–manifolds*, Princeton Math. Ser. 39, Princeton UP (1990)

[5] Bruce Kleiner and John Lott, *Notes on Perelman's papers*, arXiv:math.DG/0605667

[6] John Milnor, *Lectures on the h–cobordism theorem*, Princeton University Press (1965)

[7] John W Morgan and Gang Tian, *Ricci flow and the Poincaré Conjecture*, arXiv:math.DG/0607607

[8] Grisha Perelman, *The entropy formula for the Ricci flow and its geometric applications*, *Ricci flow with surgery on three-manifolds*, and *Finite extinction time for the solutions to the Ricci flow on certain three-manifolds*, arXiv:math.DG/0211159, arXiv:math.DG/0303109, arXiv:math.DG/0307245

[9] Henri Poincaré, *Cinquiéme complément á l'analysus situs*, Ouvres de Poincaré, Ed. Gauthier-Villars, Tome VI, Paris (1953) 435–498: Originally published: Rendiconti de Circolo Matematica di Palermo, 18 (1904) 45–110

[10] Colin Rourke and Brian Sanderson, *Introduction to piecewise-linear topology*, Springer Study Edition (1982)

[11] William P Thurston, *The geometry and topology of three-manifolds*, available on-line at: http://www.msri.org/publications/books/gt3m/

Mathematics Institute, University of Warwick
Coventry CV4 7AL, UK

cpr@msp.warwick.ac.uk