Getting Into Shapes: From Hyperbolic Geometry to Cube Complexes and Back

A proof marks the end of an era in the study of three-dimensional shapes.

By: Erica Klarreich

October 2, 2012

Thirty years ago, the mathematician William Thurston articulated a grand vision: a taxonomy of all possible finite three-dimensional shapes.

Thurston, a Fields medalist who spent much of his career at Princeton and Cornell, had an uncanny ability to imagine the unimaginable: not just the shapes that live inside our ordinary three-dimensional space, but also the far vaster menagerie of shapes that involve such complicated twists and turns that they can only fit into higher-dimensional spaces. Where other mathematicians saw inchoate masses, Thurston saw structure: symmetries, surfaces, relationships between different shapes.

Courtesy of Archives of the Mathematisches Forschungsinstitut Oberwolfach

William Thurston at Berkeley in 1991. Thurston died in August at 65.

"Many people have an impression, based on years of schooling, that mathematics is an austere and formal subject concerned with complicated and ultimately confusing rules," he wrote in 2009. "Good mathematics is quite opposite to this. Mathematics is an art of *human* understanding. … Mathematics sings when we feel it in our whole brain."

At the core of Thurston's vision was a marriage between two seemingly disparate ways of studying three-dimensional shapes: geometry, the familiar realm of angles, lengths, areas and volumes, and topology, which studies all the properties of a shape that don't depend on precise geometric measurements — the properties that remain unchanged if the shape gets stretched and distorted like Silly Putty.

To a topologist, the surface of a frying pan is equivalent to that of a table, a pencil or a soccer ball; the surface of a coffee mug is equivalent to a doughnut surface, or torus. From a topologist's point of view, the multiplicity of two-dimensional shapes — that is, surfaces — essentially boils down to a simple list of categories: sphere-like surfaces, toroidal surfaces, and surfaces like the torus but with more than one hole. (Most of us think of spheres and tori as three-dimensional, but because mathematicians think of them as hollow surfaces, they consider them two-dimensional objects, measured in terms of surface area, not volume.)

Thurston's key insight was that it is in the union of geometry and topology that three-dimensional shapes, or "three-manifolds," can be understood. Just as the topological category of "two-manifolds" containing the surfaces of a frying pan and a pencil also contains a perfect sphere, Thurston conjectured that many categories of three-manifolds contain one exemplar, a three-manifold whose geometry is so perfect, so uniform, so beautiful that, as Walter Neumann of Columbia University is fond of saying, it "rings like a bell." What's more, Thurston conjectured, shapes that don't have such an exemplar can be carved up into chunks that do.

In a 1982 paper, Thurston set forth this "geometrization conjecture" as part of a group of 23 questions about three-manifolds that offered mathematicians a road map toward a thorough understanding of three-dimensional shapes. (His list had 24 questions, but one of them, still unresolved, is more of an intriguing side alley than a main thoroughfare.)

"Thurston had this enormous talent for asking the right questions," said Vladimir Markovic, a mathematician at the California Institute of Technology. "Anyone can ask questions, but it's rare for a question to lead to insight and beauty, the way Thurston's questions always seemed to do."

These questions inspired a new generation of mathematicians, dozens of whom chose to pursue their graduate studies under Thurston's guidance. Thurston's mathematical "children" manifest his style, wrote Richard Brown of Johns Hopkins University. "They seem to see mathematics the way a child views a carnival: full of wonder and joy, fascinated with each new discovery, and simply happy to be a part of the whole scene."

In the decades after Thurston's seminal paper appeared, mathematicians followed his road map, motivated less by possible applications than by a realization that three-manifolds occupy a sweet spot in the study of shapes. Two-dimensional shapes are a bit humdrum, easy to visualize and categorize. Four-, five- and higher-dimensional shapes are essentially untamable: the range of possibilities is so enormous that mathematicians have limited their ambitions to understanding specialized subclasses of them. For three-dimensional shapes, by contrast, the structures are mysterious and mind-boggling, but ultimately knowable.

As Thurston's article approached its 30th anniversary this year, all but four of the 23 main questions had been settled, including the geometrization conjecture, which the Russian mathematician Grigori Perelman proved in 2002 in one of the signal achievements of modern mathematics. The four unsolved problems, however, stubbornly resisted proof.

"The fact that we couldn't solve them for so long meant that something deep was going on," said Yair Minsky, of Yale University.

Finally, in March, Ian Agol, of the University of California at Berkeley, electrified the mathematics community by announcing a proof to "Wise's conjecture," which settled the last four of Thurston's questions in one stroke.

Mathematicians are calling the result the end of an era.

"The vision of three-manifolds that Thurston articulated in his paper, which must have looked quite fantastic at the time, has now been completely realized," said Danny Calegari, of the California Institute of Technology. "His vision has been remarkably vindicated in every way: every detail has turned out to be correct."

"I used to feel that there was certain knowledge and certain ways of thinking that were unique to me," Thurston wrote when he won a Steele mathematics prize this year, just months before he died in August at 65. "It is very satisfying to have arrived at a stage where this is no longer true — lots of people have picked up on my ways of thought, and many people have proven theorems that I once tried and failed to prove."

Agol's result means that there is a simple recipe for constructing all compact, hyperbolic three-manifolds — the one type of three-dimensional shape that had not yet been fully explicated.

"In a precise sense, we now understand what all three-manifolds look like," said Henry Wilton, of University College London. "This is the culmination of a massive success story in mathematics."

Surface Study

Thurston's program tried to do for three-dimensional manifolds what mathematicians had successfully done more than a century earlier for two-dimensional manifolds. As a warm-up for understanding three-dimensional manifolds, let's look under the hood at the classification of "compact, orientable" surfaces (finite surfaces with no punctures or gashes and a consistent sense of orientation).

To tackle this classification problem, mathematicians showed that, given an arbitrary surface, it is possible to progressively simplify it by cutting it open along curves until the surface completely opens out into a flat polygon.

It's easy to see how to do this with, say, a torus: first cut it open along loop A in Figure 1, producing a cylinder. Next, cut along loop B, flattening the cylinder out into a square. It's a little harder to see, but cutting along the four curves in Figure 2 converts a double torus (a torus with two holes) into an octagon. Similarly, for any n-holed torus, we can cut along $2n$ loops to flatten out the surface into a $4n$-gon.

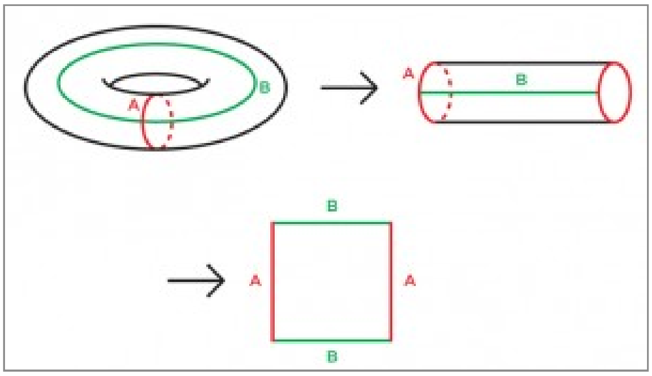

Figure 1. Cutting a torus open along loop A yields a cylinder. Cutting further, along loop B, unfurls the cylinder into a square.

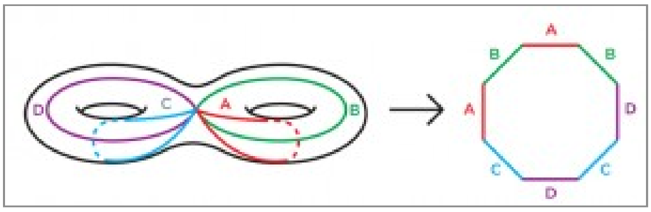

Figure 2. Cutting a double torus along loops A, B, C and D yields an octagon.

Given an arbitrary, unidentified surface, we can try to simplify it (and ultimately identify it) by dissecting it in a similar way. Provided that the surface is not a sphere, topologists have shown that it must contain some embedded loops (loops that don't intersect themselves) that cannot be pulled down to a single point, similar to loops A and B on the torus. Dissecting the surface along one of these loops removes some of the surface's interesting topological features. Mathematicians have shown that there are only a finite number of times we can cut in this way before we have reduced the surface to a flat polygon.

Once we have simplified our surface down to a polygon, it's fairly simple to see that when we re-glue the sides to recover our original surface, we must produce a torus, or a double torus, or a triple torus, and so on. After all, the first gluing will turn the polygon into a tunnel-shaped surface, and then each subsequent gluing will either introduce a new tunnel-shaped handle on the surface or simply sew up some open seams. When we're finished, the result is a torus surface with some number of holes.

This approach does more than just show that the surface is topologically equivalent to a sphere or a torus of some type: it also gives a way to endow the surface with a simple, uniform geometric structure.

A sphere clearly already has a uniform geometric structure: its geometry looks the same no matter where you're standing on the surface. A doughnut surface, by contrast, is anything but uniform: a region on the outer edge of the doughnut curves in a way that's reminiscent of a sphere, while a region on the inner ring of the doughnut curves more like the surface of a saddle.

No matter how you try to place a torus in space — no matter how much stretching and distorting you do — there's no way to make its geometry look the same at every point. Some parts will curve like a sphere and some like a saddle, and some parts may be flat.

Nevertheless, it's possible to equip the torus with an abstract geometric structure that is identical at every point: simply declare that on each small patch of the torus, distances and angles are to be measured by taking the corresponding measurements on the square from which, as we've seen, the torus can be built. It's impossible to build a physical torus inside ordinary space whose lengths and angles match this abstract rule, but this definition of lengths and angles is internally consistent. Since the square has ordinary flat (Euclidean) geometry, we say that the torus can be equipped with a Euclidean structure. A torus with this geometry is akin to a video game in which, when a creature exits the screen on the right, it reappears on the left, and when it exits at the top of the screen, it reappears at the bottom.

If we try to do the same thing for the double torus, however, we hit a snag. Recall that we can build a double torus by gluing the edges of an octagon. If we declare that geometry on the double torus shall mimic geometry on the octagon, we run into a problem at the octagon's corners. After the octagon has been glued up into a double torus, the corner points are all glued together to form a single point on the double torus. Eight octagon corners meet up at that point, each corner contributing 135 degrees of angle measure, for a total of 1080 degrees, instead of the usual 360 degrees.

So if we try to give the double torus the same geometric structure as the octagon, we will end up with a double torus that has ordinary Euclidean geometry everywhere except at one point, where the surface buckles like a floppy hat with a sharp peak. (The corner points are not a problem when we glue a square to make a torus: we glue four 90-degree corners to get a perfect 360 degrees.)

To get a smooth geometric structure at the corner point on the double torus, we would need each of the octagon's eight corners to contribute 45 degrees instead of 135 degrees. Remarkably, such an octagon does exist, but it lives not in the ordinary Euclidean plane but in another geometric structure called the hyperbolic disk: a third kind of geometry which is as uniform and internally consistent as spherical or Euclidean geometry, but which, because it is harder to visualize, was not even discovered by mathematicians until the early 19th century.

Roughly speaking, hyperbolic geometry is what you get if you declare that all the fish in Figure 3 are the same size. It's as if Figure 3 is really the image of the hyperbolic disk through a distorted lens that makes the fish near the boundary look much smaller than the fish in the middle. In the real hyperbolic disk that is theoretically on the other side of the lens, the fish are all identical in size.

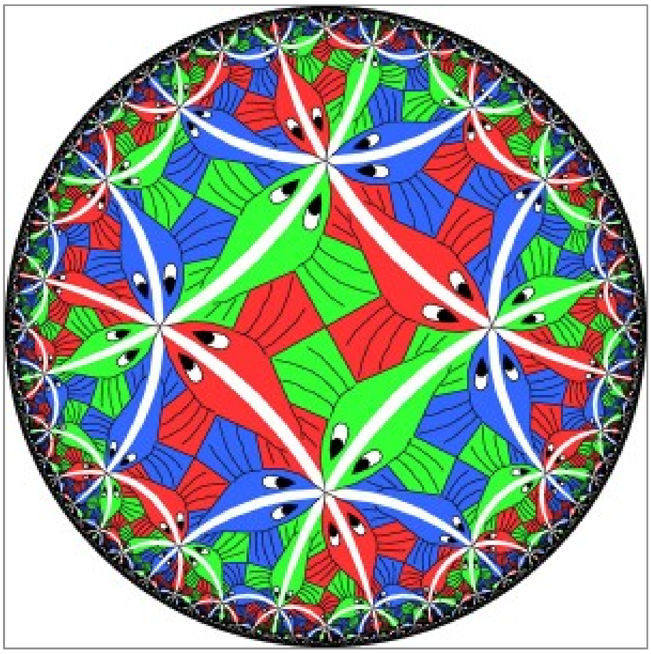

Illustration by Douglas Dunham, University of Minnesota Duluth

Figure 3. When viewed through the lens of hyperbolic geometry, all the fish are the same size. The curves that run along the fishes' spines are hyperbolic straight lines, or "geodesics."

There's no way to make a nice, smooth hyperbolic disk in ordinary space so that the fish truly are the same size. But once again, from an abstract point of view, the fish-sizing rule produces a geometry that is internally consistent and looks the same at every point — not when viewed by an outsider looking through the distorted lens, but from the perspective of someone who lives in the hyperbolic disk.

In hyperbolic geometry, the shortest path, or "geodesic," between two points is the path that travels through the fewest possible fishes to get from one point to the other. Such a path, it turns out, is always a semicircle perpendicular to the boundary of the disk; the semicircles that go through the fishes' spines are examples. From our distorted outside perspective, such paths look curved, but for an insider, these paths are the "straight lines": to drive along one of them, you would never have to turn the steering wheel, as Thurston often put it. In contrast with the Euclidean plane, in which parallel lines always stay the same distance apart, in the hyperbolic disk, two lines that don't intersect can spread apart from each other very quickly.

From the point of view of hyperbolic geometry, the shapes in Figure 4 are all regular octagons with straight edges. In one of these octagons, the angles are all 45 degrees — just what we need for a double torus. If we glue this octagon's sides appropriately, the result will be a double torus with a perfect, uniform hyperbolic structure.

Similarly, we can equip a triple torus with a hyperbolic structure. A triple torus can be made by gluing the sides of a 12-sided polygon, so if we construct a hyperbolic dodecagon whose internal angles are all 30 degrees, its hyperbolic geometry can be carried over smoothly to the triple torus. Continuing in this way, we can equip a four-holed torus, a five-holed torus, and so on, with hyperbolic geometry. Our taxonomy of compact surfaces becomes: one surface with spherical geometry (the sphere), one surface with Euclidean geometry (the torus), and infinitely many surfaces with hyperbolic geometry (all the tori with more than one hole).

Over the past century, this taxonomy has given mathematicians an incredibly fruitful way to translate topological questions about surfaces into geometric ones, and vice versa. The classification of surfaces is the foundational concept in the study of two-dimensional shapes, a finding that all subsequent studies take as their starting point.

The Next Dimension

Three-manifolds are vastly more diverse than two-manifolds, and the questions are correspondingly harder. Even as simple-sounding a question as the famous Poincaré conjecture — which asks whether the three-dimensional version of the sphere is the only compact three-dimensional shape on which every loop can be pulled tight to a single point without getting snagged around a hole — remained unsolved for nearly a century after Henri Poincaré posed it in 1904.

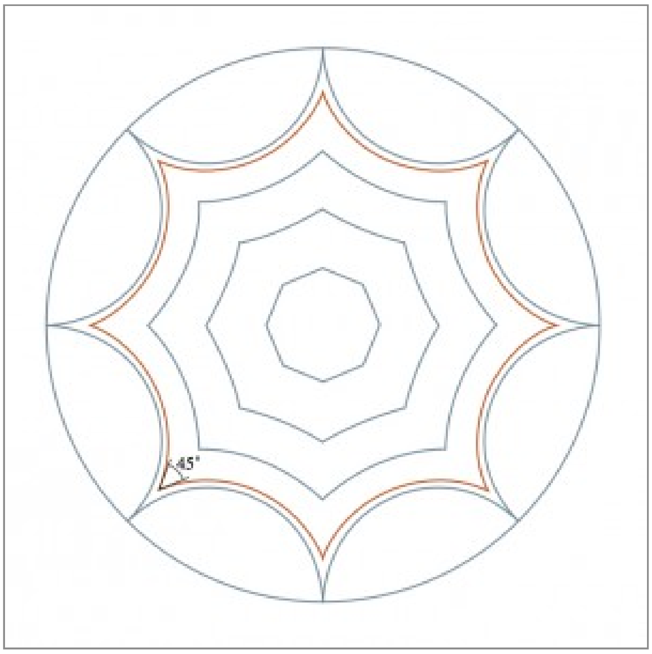

Illustration by Silvio Levy)

Figure 4. Regular octagons in hyperbolic space, such as the ones pictured above, can have any internal angle measure greater than zero and less than 135 degrees. The brown octagon, whose internal angles are all 45 degrees, can be glued together to form a double torus with smooth hyperbolic geometry.

Nevertheless, Thurston boldly conjectured that it should be possible to create a taxonomy for three-dimensional shapes similar to the one for two-dimensional shapes.

The two-dimensional Euclidean, spherical and hyperbolic geometries each have a counterpart in three dimensions. But in three dimensions, these are not the only "nice" geometries out there. For example, there are hybrid geometries that are hyperbolic or spherical in certain directions, and Euclidean in others. Altogether, there are eight different types of geometry in dimension three that are uniform, meaning that the geometry looks the same at every point in the space.

Thurston conjectured that, just as with surfaces, three-manifolds can be endowed with natural geometric structures. Specifically, he proposed that if you carve up any compact three-manifold into chunks in a particular way, each chunk can be endowed with one of the eight geometries.

"The goal was to completely unify topology and geometry in three dimensions," Minsky said.

A natural approach to proving this "geometrization conjecture" would be to try to do something similar to what we did for surfaces, which we cut along curves until we had cut open all the interesting topological features and reduced the surface to a flat polygon. For a three-manifold, the corresponding approach would be to cut it open along surfaces until, hopefully, it reduced to a polyhedron, whose opposite sides could be glued together to recover the original shape. Then, if we could build the polyhedron with the right geometry, we could transfer that geometry to the original shape, just as we did with surfaces.

Remember that to make this work for surfaces, each curve we cut along must satisfy two properties: The curve should never cross itself (in mathematical lingo, it should be "embedded"), and it should be what we'll call "topologically interesting," meaning that it winds around some topological feature of the surface and can't be tightened down to a single point (this requirement ensures that cutting along the curve simplifies the surface topologically).

In 1962, the mathematician Wolfgang Haken proved that it is indeed possible to simplify a three-manifold down to a polyhedron, provided the three-manifold contains a surface to cut along that satisfies two conditions: It must be embedded, and it must be "incompressible," meaning that every topologically interesting curve on the surface is also topologically interesting in the larger context of the surrounding three-manifold.

So, for example, a torus is not incompressible in ordinary three-dimensional space, since the loop that dips through the hole of the torus is topologically interesting from the point of view of the torus, but in the full three-dimensional space it can be compressed down to a single point. By contrast, the torus is incompressible inside the three-manifold that you get just by thickening the torus surface slightly so that it is no longer infinitesimally thin. To be incompressible, every topological feature of the surface must truly reflect some of the three-manifold's intrinsic topology. A three-manifold that has an embedded, incompressible surface is now known as a Haken manifold.

If our three-manifold does have an embedded, incompressible surface, then cutting along this surface will cut open some of the three-manifold's interesting topology, leaving a simpler manifold in its place. What's more, Haken showed that as long as the manifold contains one such surface, the new manifold produced by the cutting will itself be Haken: it will again have an embedded, incompressible surface to cut along. After a finite number of such steps, Haken showed, all the interesting topological features of the original manifold will have been cut away, leaving a polyhedron.

In the late 1970s, Thurston showed that it is possible to endow the resulting polyhedron with one of the eight three-dimensional geometries in such a way that the geometry transfers smoothly to the re-glued polyhedron, fitting together perfectly at the polyhedron's corners and edges. In other words, Thurston proved his geometrization conjecture for manifolds whose standard decomposition yields chunks that are all Haken manifolds.

Unfortunately, given an arbitrary compact three-manifold, there's no guarantee that it will indeed have such a surface. In fact, in the late 1970s and early 1980s, Thurston convinced the three-manifold community that three-manifolds that contain an embedded, incompressible surface (Haken manifolds) are the exception, rather than the rule.

Figuring out how to prove the geometrization conjecture for non-Haken manifolds stumped mathematicians for more than two decades. Finally, in 2002, Perelman set forth his proof, which drew on areas of mathematics far removed from those studied by most of Thurston's followers. (Along the way, Perelman's proof settled the century-old Poincaré conjecture, leading the Clay Mathematics Institute in 2010 to offer him a million-dollar prize — which he promptly rejected, for rather complicated reasons.)

Perelman's landmark proof achieved Thurston's dream of unifying topology and geometry. Now every topological question about three-manifolds had its geometric counterpart, and vice versa. But Perelman's theorem left unresolved many important questions about what kinds of three-manifolds can exist.

In classifying compact two-manifolds (surfaces), mathematicians were able not only to show that each surface could be endowed with a geometric structure, but also to make a complete list of every possible two-manifold. In dimension three, such a list was lacking.

Seven of the eight three-dimensional geometries — all but hyperbolic geometry — are fairly easily understood, and even before Perelman's work, three-manifold topologists had arrived at a complete description of the types of manifolds that can admit one of these seven geometries. Such shapes are relatively simple and few.

But just as with surfaces, in dimension three it turns out that most manifolds are in fact hyperbolic. Mathematicians had a much more tenuous grasp of the vast range of possibilities for hyperbolic three-manifolds than they had for the other seven geometries.

"Of the eight kinds of geometry, the hyperbolic manifolds are the most mysterious and rich," said Nicolas Bergeron, of the Université Pierre et Marie Curie in Paris.

Perelman's result told mathematicians that hyperbolic manifolds were indeed the final frontier — the only kind of three-manifold left to understand. But it didn't tell them what these hyperbolic shapes actually look like.

Cover Story

Once again, mathematicians were able to turn to Thurston's seminal paper for guidance. On his famous list of questions were many conjectures about the features of hyperbolic three-manifolds, including two conjectures that speak directly to what such manifolds can look like: the "virtual Haken" conjecture and the "virtual fibering" conjecture.

The virtual Haken conjecture proposes that every compact hyperbolic three-manifold is almost Haken, in a precise sense: it's possible to convert the manifold into a Haken manifold simply by unrolling it a finite number of times, in a particular way. This new, unrolled manifold is called a "finite cover" of the original manifold.

Mathematicians say that one manifold N covers another manifold M if, roughly speaking, it's possible to wrap N around M a certain number of times (perhaps infinitely many times) so that each part of M gets covered the same number of times as every other part. To be a covering, this wrapping should also have an assortment of other nice properties — for example, N should never fold over on itself or tear during this wrapping process. Each little piece of M is covered by a bunch of identical copies of it in the cover, N.

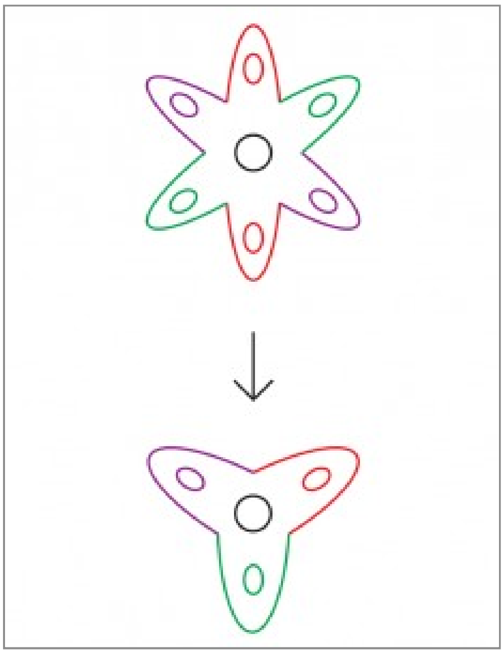

Figure 5. The six-petal flower covers the three-petal flower by wrapping around it two times.

For example, the six-petal flower in Figure 5 covers the three-petal flower: simply wrap the six-petal flower two times around the three-petal flower. Each point on the three-petal flower is covered by two points on the six-petal flower; mathematicians call this a two-sheeted covering.

Likewise, an infinitely long cylinder covers a torus: just keep wrapping the cylinder evenly around and around the torus, infinitely many times (see Figure 6). Every point on the torus is covered: Loop A on the torus is covered by an infinite collection of evenly spaced loops on the cylinder, while loop B unrolls on the cylinder to become a line that runs the length of the cylinder.

The topology of a manifold and its cover are intimately related. To reconstruct a manifold from an n-sheeted cover, you simply fold the cover over on itself n times. Likewise, to reconstruct the cover from the manifold, you slice open the manifold, make n copies of it, and glue the copies together along their boundaries (the particular cover you get may depend on your gluing choices).

A cover preserves some of the manifold's topological features while unrolling others. The infinite cylinder, for example, remembers that loop A is a closed loop in the torus, but it forgets that loop B is also a closed loop.

This unrolling process is precisely what led Thurston to hope that, given a three-manifold, it might be possible to produce a finite-sheeted cover that is Haken. As we've discussed, given an arbitrary compact, hyperbolic three-manifold, there is no reason to expect it to be Haken (that is, to have an embedded, incompressible surface). However, in 1968, German mathematician Friedhelm Waldhausen conjectured that such a manifold should at least contain an incompressible surface, although the surface might pass through itself in places, rather than being embedded.

If that is indeed the case, Thurston argued, there might well be a finite cover in which the surface unrolls in a way that eliminates all of its intersections with itself. Finite covers can often achieve such simplifications. For example, since the curve in the three-petal flower in Figure 7 goes around the central hole twice, no amount of stretching and shifting can prevent it from intersecting itself somewhere. But if we start unrolling this curve in the six-petal flower starting at a chosen point P, the resulting red curve (which mathematicians call a "lift" of the original curve) goes around the central hole only once and never intersects itself. (There is a second lift, the blue curve, which intersects the red curve at the two points that cover the intersection point in the three-petal flower.)

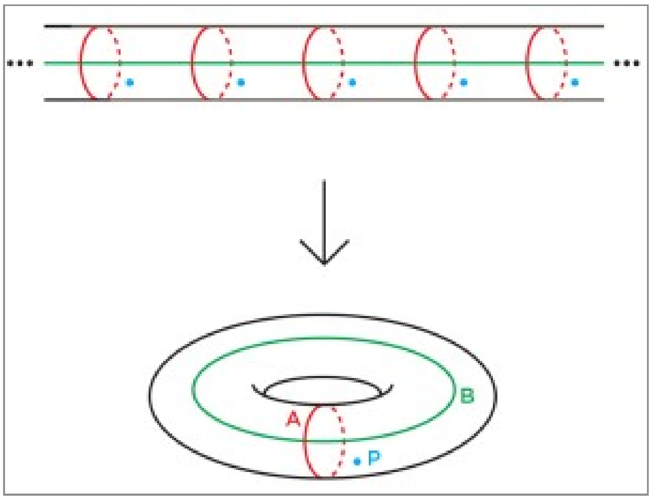

Figure 6. An infinitely long cylinder covers a torus by wrapping around it again and again. Loop A on the torus "lifts" to the infinite collection of red loops on the cylinder. Loop B unrolls in the cylinder to become the green line. Point P on the torus lifts to the infinite collection of blue dots on the cylinder.

In his 1982 paper, Thurston proposed that given a compact, hyperbolic three-manifold, it should be possible to do a similar type of unrolling to produce embedded surfaces in some finite cover — in other words, the three-manifold should be "virtually" Haken.

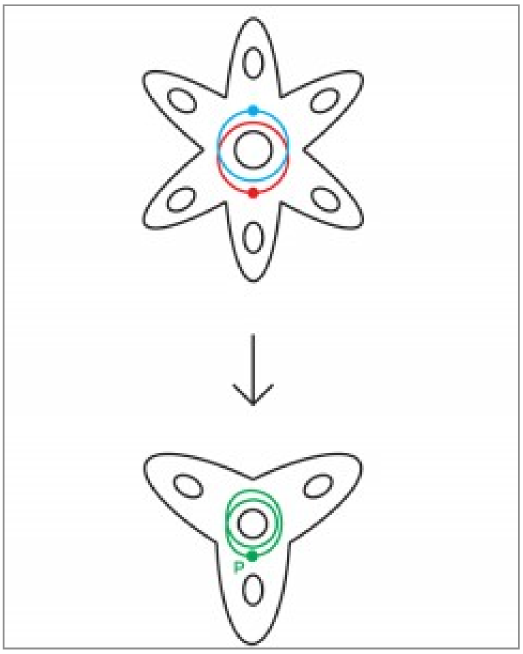

Figure 7. The green curve in the three-petal flower intersects itself, but its two lifts in the six-petal flower, the red and blue curves, never intersect themselves (though they intersect each other).

A Haken manifold, as we've discussed, can be built by gluing the boundary walls of a polyhedron in a particular way. The virtual Haken conjecture implies, then, that any compact hyperbolic three-manifold can be built first by gluing up a polyhedron nicely, then by wrapping the resulting shape around itself a finite number of times.

Thurston went on to suggest something even stronger: that any compact hyperbolic three-manifold might be virtually fibered, meaning that it has a finite cover that is "fibered." A manifold that "fibers over the circle" (as mathematicians say) is built by thickening a surface slightly to make it three-dimensional, then gluing the inner and outer boundary surfaces together according to any arrangement that matches the two surfaces up smoothly, point for point. (Such a gluing couldn't be realized in ordinary space without parts of the resulting manifold passing through itself, but it can still be studied abstractly.) The manifold is said to be fibered because if you imagine stretching out the thickened surface so the boundary surfaces are far apart, then drawing the boundaries around to face each other before gluing them together, you can imagine that the resulting manifold is like a bracelet that has an infinitely thin surface-shaped bead at every point on the bracelet's strand; these beads are the "fibers."

Every fibered manifold is Haken, but the reverse is not true. Thus, the virtual fibering conjecture is a stronger statement than the virtual Haken conjecture, and Thurston was on the fence about whether it was indeed true. "This dubious-sounding question seems to have a definite chance for a positive answer," was as far as he was willing to go in his 1982 paper.

Thurston had originally proposed the virtual Haken conjecture in an early attempt to tackle his geometrization conjecture, which he had already proven for Haken three-manifolds. If the virtual Haken conjecture were true, so that every compact three-manifold has a Haken finite cover, it might be possible, Thurston hoped, to use the geometric structure on the cover to build a geometric structure on the original manifold.

Three decades later, well after Perelman proved the geometrization conjecture by very different means, the virtual Haken conjecture and the virtual fibering conjecture remained unsolved. These, and two other related conjectures, were the only questions left unanswered among Thurston's 23. Computer data strongly suggested that the virtual Haken conjecture was correct: from a computerized list of more than 10,000 hyperbolic three-manifolds, Thurston and Nathan Dunfield, of the University of Illinois at Urbana-Champaign, had managed to find a Haken finite cover for every single one. But computational evidence is not a proof.

"When Thurston proposed it, the virtual Haken conjecture seemed like a small question, but it hung on stubbornly, shining a spotlight on how little we knew about the field," Minsky said. "It turned out that our ignorance was deep in that direction."

Building Surfaces

In 2009, the murky waters surrounding the virtual Haken conjecture started to clear.

That year, Markovic and Jeremy Kahn, then at Stony Brook University and now at Brown, announced the proof of a key step toward proving the virtual Haken conjecture. The result, which we'll call the "incompressible surface theorem," states that every compact hyperbolic three-manifold does indeed contain an incompressible surface (which possibly passes through itself instead of being embedded).

Kahn and Markovic's proof is a prime example of the interplay between three-dimensional topology and geometry: the incompressible surface theorem is a purely topological statement, but to prove it, Kahn and Markovic drew heavily on the wealth of additional structure that hyperbolic geometry provides.

To build surfaces inside a three-manifold, Kahn and Markovic used an attribute of hyperbolic shapes called "exponential mixing." This means that if you start in any little neighborhood inside your manifold, pick a direction, and imagine that your neighborhood is starting to flow along a river moving roughly in that direction, then your neighborhood will gradually spread out and wind around the three-manifold, reaching every possible location from every possible direction. What's more, it will spread out very quickly, in a precise "exponential" sense.

This mixing property is unique to hyperbolic three-manifolds and stems roughly from the fact that, unlike in Euclidean space, in hyperbolic space the "straight lines," or geodesics, curve away from each other. If you pick a small neighborhood in the hyperbolic disk and let it flow in a particular direction, the neighborhood will grow exponentially quickly. Inside a compact three-manifold, a flowing neighborhood will likewise grow exponentially quickly, but since the entire manifold has a finite extent, the neighborhood will end up winding around the manifold again and again, overlapping itself many times. Furthermore — and this is harder to prove — the neighborhood will wind around the manifold evenly, flowing through all spots in the manifold with roughly the same frequency.

Mathematicians have understood this exponential-mixing property for more than 25 years and have thoroughly analyzed the statistics of this "geodesic flow," figuring out roughly when and how often a given neighborhood will pass by a particular point as the neighborhood flows along. But until Kahn and Markovic tackled the incompressible surface theorem, mathematicians had never successfully harnessed this mixing property in the service of building topological structures in a manifold. (One other mathematician, Lewis Bowen of Texas A&M University, had previously tried to use exponential mixing to build incompressible surfaces in three-manifolds, but his work hit technical obstacles.)

To see how the exponential-mixing property helps to build topological and geometric structures, let's apply it to a simpler task than building surfaces: building a closed geodesic loop whose length is close to our favorite large number — call it R.

To build our loop, let's pick any starting point in the manifold and any starting direction, and then imagine turning on a garden hose located in a little neighborhood surrounding that point, and aimed roughly in that direction. The water droplets will stream out along geodesic paths, and as long as R is sufficiently large, the mixing of the flow means that by the time the droplets have traveled a distance R, they will have spread out fairly evenly throughout the whole manifold. In particular, at least one droplet (in fact, many) will have arrived back near the starting point and the starting direction. Then, we can simply build a little bridge connecting that droplet's geodesic to the starting point, to produce a loop that is almost perfectly geodesic and whose length is very roughly equal to R. It's not hard to show that by pulling this loop just a bit tighter in the manifold, we can produce a totally geodesic loop.

Notice that this method didn't give us just one geodesic loop of length close to R. Any starting point and starting direction can be used in this process, and many of the water droplets will come back near the starting point, so in fact we can generate many such loops. This is a general principle of structure building using exponential mixing.

Exponential mixing "says that whatever structures you find in your manifold, you'll find in abundance," Calegari said.

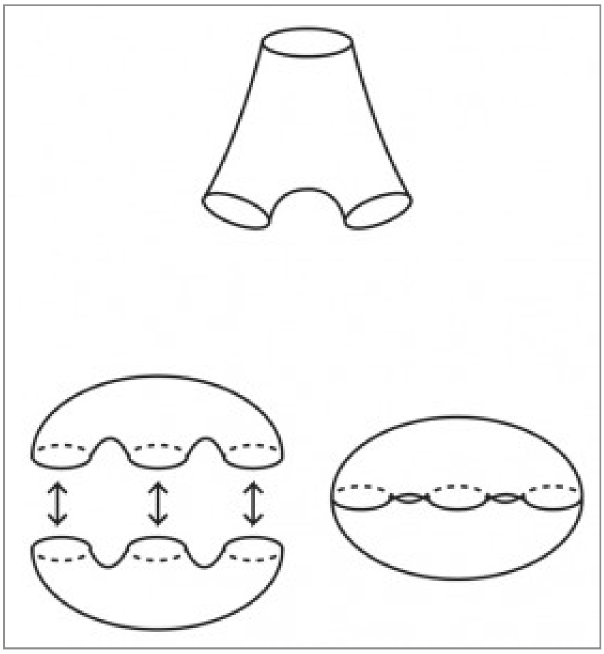

Figure 8. A pair of pants (top); gluing together two pairs of pants (bottom left) produces a double torus (bottom right).

Kahn and Markovic used a similar approach to our loop-building exercise to build "pairs of pants" — surfaces topologically equivalent to a sphere with three holes (a waist hole and two leg holes, so to speak). Pairs of pants are the building blocks of all compact surfaces except the sphere and the torus — for example, gluing together two pairs of pants produces a double torus (see Figure 8).

Given any sufficiently large number R, Kahn and Markovic showed that it is possible to build lots of pairs of pants inside the manifold whose three cuffs each have a length close to R, and that are almost totally geodesic, meaning that each bit of the pants surface looks pretty much flat from the point of view of hyperbolic geometry.

They also showed that at each cuff of a pair of pants, there is another pair of pants emanating from the cuff in roughly the opposite direction. By sewing together these matching pants at the cuffs, Kahn and Markovic produced a large family of compact surfaces that are almost totally geodesic, with some slight buckling at the seams. Surfaces that are almost geodesic are known to be incompressible inside their three-manifold, so Kahn and Markovic's construction proved the incompressible surface theorem.

Their methods also showed that a three-manifold has not just one incompressible surface, but "a rich structure of almost geodesic surfaces all over the place," Calegari said.

Kahn and Markovic's work earned them the 2012 Clay Research Award, presented annually by the Clay Mathematics Institute to recognize major mathematical breakthroughs.

"The techniques of Kahn and Markovic are as compelling as their results, and this body of work will undoubtedly inspire many more threads of inquiry than it ties off," predicted Jeffrey Brock of Brown University, in an article on Kahn and Markovic's work in late 2011.

A Hidden Structure

For mathematicians trying to prove the virtual Haken conjecture, Kahn and Markovic's work created a starting point.

They showed that every manifold is guaranteed to contain an incompressible surface. But this surface may pass through itself, perhaps in many places, instead of being embedded. To get from Kahn and Markovic's result to the virtual Haken conjecture, mathematicians would have to find a finite cover of the manifold in which, just as in the example of the six-petal and three-petal flowers, the surface lifts to a collection of surfaces that never intersect themselves (though they may intersect each other). If this could be done, each of these would be an embedded, incompressible surface in the cover, meaning the cover would be Haken.

But how, exactly, is such a cover to be found?

"There's a big gap between Kahn and Markovic's result and the virtual Haken conjecture," Dunfield said. "Their finding was important, but at the time it wasn't so clear whether it would be helpful in getting embedded surfaces."

Kahn and Markovic's result caught the attention of Daniel Wise, of McGill University in Montreal. Wise had, in a sense, made a career of figuring out when finite covers remove a topological object's self-intersections, but he worked in the context of "cube complexes," objects that are seemingly very different from three-manifolds. Kahn and Markovic's findings allowed Wise to show other mathematicians that these two contexts are not so very far apart.

A cube complex is just what it sounds like: a collection of cubes, except that the word "cube" refers not just to the usual three-dimensional cube but to the shape in any dimension consisting of all the points whose coordinates lie between, say, -1 and +1. For example, a square is considered a two-dimensional cube, and a line segment is a one-dimensional cube. The cubes in a cube complex are connected to each other along corners, edges, faces and higher-dimensional sides.

Cube complexes are very different creatures from three-manifolds — they aren't even manifolds, for starters, since the junctions between two cubes of different dimensions don't resemble ordinary space of any dimension. Yet cube complexes are a simplified setting in which to study one key aspect of a surface sitting inside a three-manifold: the fact that such

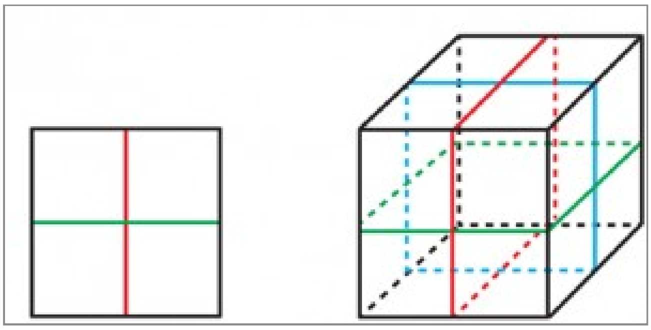

Figure 9. A square (left) has two hyperplanes (the red and green lines). A cube has three hyperplanes (the red, blue and green squares).

a surface, at least locally, divides its surroundings into two sides.

If your goal is to study objects that divide a shape into two sides, cubes are a natural place to start, since of all possible shapes, they have some of the simplest such objects: the "hyperplanes" that cut across the middle of the cube. A square has two hyperplanes — the vertical and horizontal lines that each chop the square in half — and a cube has three hyperplanes (see Figure 9). An n-dimensional cube has n hyperplanes, which all intersect at the center point of the cube.

"The hyperplanes are like surfaces in a three-manifold, but you see them immediately," Wise said. "Finding surfaces is hard, but hyperplanes are available to you right to begin with."

If we start with a hyperplane inside a cube in a cube complex, there is exactly one way to extend the hyperplane to hyperplanes in the adjacent cubes; after that, there's exactly one way to extend *those* hyperplanes to *their* adjacent cubes; and so on. Thus, given a starting hyperplane in a cube complex, there's a unique way to extend it to a hyperplane in the full cube complex (see Figure 10).

This quality provides a stark contrast to three-manifolds, in which a small piece of surface can be extended in any number of ways to a full surface. Cube complexes and their hyperplanes are "nice and crystalline and rigid,"Agol said, with none of the "flabbiness" of a three-manifold and its surfaces.

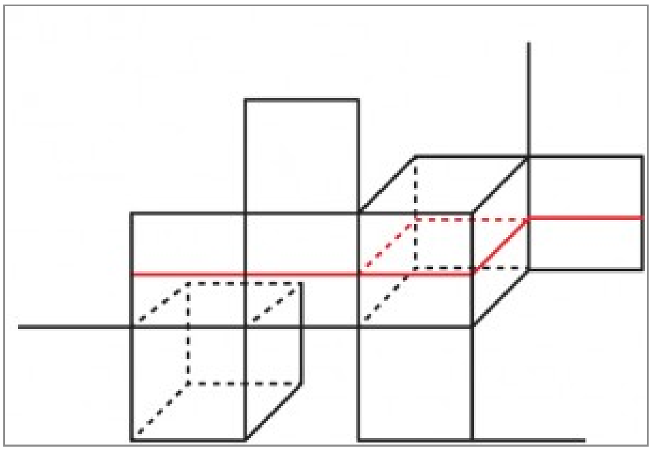

Figure 10. The red hyperplane in the rightmost square extends in a unique way to a hyperplane in the cube complex as a whole.

As we extend a hyperplane through a cube complex, it may come back to the cube where it started and pass through it along a hyperplane perpendicular to the original one (see Figure 11). In other words, the extended hyperplane might not be embedded. Just as with surfaces inside three-manifolds, we can ask whether the cube complex has a finite cover in which these self-intersecting hyperplanes lift to embedded ones — the cube-complex version of being virtually Haken.

Several years ago, Wise and Frédéric Haglund, of the Université Paris-Sud in Orsay, France, defined a class of "special" cube complexes which, among other nice features, have only

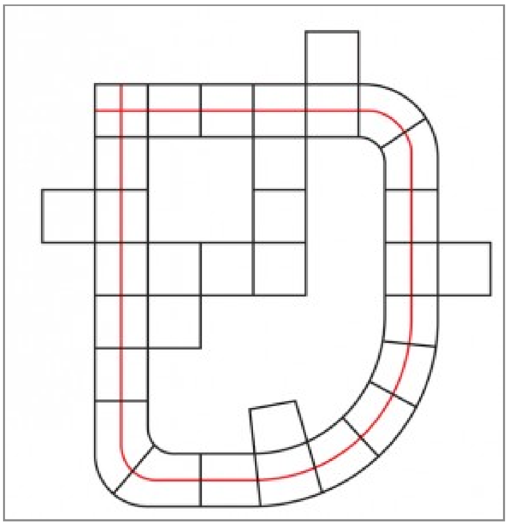

Figure 11. As we extend the horizontal red hyperplane in the square at upper left through other parts of the cube complex, it comes back around and intersects itself perpendicularly.

embedded hyperplanes. Over the course of the last decade, Wise has developed an arsenal of techniques for figuring out which cube complexes are special. In 2009, Wise circulated a 200-page "masterwork," as Dunfield put it, in which he detailed a host of findings about special cube complexes, such as his "combination theorems," which show how to piece together special cube complexes to get new ones that are guaranteed to still be virtually special. In this paper, Wise formulated a conjecture that stated, very roughly, that any cube complex whose geometry curves around in a similar way to hyperbolic geometry is "virtually" special — that is, it has a special finite cover. This statement came to be known as Wise's conjecture.

Wise was convinced that when a given shape is akin to a cube complex in a particular way — that is, when it can be "cubulated" — the structure of the cube complex is the key to unlocking many attributes of the original shape.

"The cube complex was a secret that people didn't even know to ask about," he said. "It's a fundamental, hidden intrinsic structure."

Cubical Scaffolding

Wise became "crazy excited" about cubulating shapes, he said, but at first his mathematical friends just laughed at his monomania.

Then Kahn and Markovic proved the incompressible surface theorem, and Wise and Bergeron immediately published a paper showing that the existence of incompressible surfaces in a compact hyperbolic three-manifold gave a way to cubulate it — and in such a way that the surfaces in the three-manifold correspond precisely to the hyperplanes in the resulting cube complex.

The key to Wise and Bergeron's construction was the fact that Kahn and Markovic had shown how to construct not one, but a multitude of surfaces. Following an approach to cubulation pioneered in 2003 by Michah Sageev, now at the Technion in Haifa, Israel, Wise and Bergeron started by taking a huge collection of Kahn-Markovic surfaces — enough to divide the three-manifold into compact polyhedra.

Now consider one of the intersection points of these surfaces — suppose, say, that *n* surfaces meet at this point. Sageev's insight was to view such an intersection as the shadow, so to speak, of the intersection of n hyperplanes in an n-dimensional cube. The cube complex corresponding to the three-manifold is built, roughly, by putting in one n-dimensional cube for each intersection of n surfaces (the actual construction is a bit subtler, in order to deal with various topological contingencies). Two cubes in the complex are adjacent if their corresponding intersection points in the three-manifold are connected by a face of one of the polyhedra.

"The cube complex is there precisely to record how the surfaces intersect with themselves and each other," Dunfield said.

Wise and Bergeron showed that this cube complex is "homotopy equivalent" to the original manifold, meaning that the cube complex can be squished and stretched around (perhaps with some dimensional flattening and unflattening) until the cube complex has turned into the manifold, and vice versa. What's more, this homotopy equivalence transforms each surface in the three-manifold into a corresponding, homotopy-equivalent hyperplane in the cube complex.

The cube complex constructed in this way satisfies the geometric requirements of Wise's conjecture, meaning that if Wise's conjecture is true then this cube complex has a finite cover in which all the hyperplanes are embedded.

If such a finite cover indeed exists (let's say, a cover with m sheets), then recall that the cover could be built from the cube complex by cutting the complex open in some way, making m copies of the complex, and gluing the copies together along the cuts. It's not hard to show that this recipe for making the cover would carry directly over to a corresponding recipe for making a finite cover of the three-manifold, and that in this finite cover, the Kahn-Markovic surfaces that were used to build the cube complex would lift to embedded surfaces. In other words, if Wise's conjecture is true, then so is the virtual Haken conjecture.

"The tradeoff is strange: your cube complex might be 10,000-dimensional, for example, so on some level it seems as if you're making things worse," Wise said. "But even though the cube complex is so big, many features about it are very easy to understand, so it's very valuable. We prefer to have something that is big but well-organized rather than to have a three-manifold."

Even after Wise and Bergeron made the connection between cube complexes and the virtual Haken conjecture, most three-manifold topologists kept their distance from cube complexes. Perhaps this was because Wise's 200-page paper seemed so daunting, or because cube complexes were so different from the kinds of spaces they were used to studying.

"These ideas were quite esoteric for people coming from hyperbolic geometry," Bergeron said.

But one mathematician was already fluent in both three-manifold topology and the more abstract, combinatorial considerations that were the currency of Wise's approach.

"I think Ian Agol was the only three-manifold guy who understood very early on that Wise's ideas were useful for three-manifold topology," Bergeron said.

Agol dived into the study of Wise's masterwork and became convinced that all the parts of it pertaining to Wise's conjecture were indeed correct. Agol had been working for some time on the virtual Haken conjecture; he realized that Wise's approach, which resolved the flabby surfaces into crystalline hyperplanes, was exactly what he needed.

"The cube complex gives a scaffolding on which to construct the finite cover," he said.

To build a special finite cover of a Wise-Bergeron cube complex, Agol started by (abstractly) chopping up the cube complex into "Lego blocks," cutting along the hyperplanes. He then assigned colors to the faces of the blocks so that any two faces that meet at a corner have different colors. Next, Agol showed, roughly speaking, that there is a way to glue together a finite number of copies of the Lego blocks along faces with matching colors, in such a way that the colors on the sides of those faces also match; that way, each extended hyperplane will be all one color. The resulting cube complex is a finite cover of the original one, and all of its hyperplanes are embedded, since any two hyperplanes that intersect have different colors and therefore are not the same hyperplane intersecting itself.

On March 12, Agol announced that he had proven Wise's conjecture, and thus the virtual Haken conjecture.

"It was the most exciting news since Perelman proved the geometrization conjecture," Dunfield said.

Word raced through the three-manifold community, and cube complexes suddenly became a common topic of conversation among three-manifold topologists.

"Until now, I don't think the mathematics community had realized just how powerful Wise's work is," Agol said. "I think my result will make people more aware of what spectacular progress he has made."

Now, Wise said, mathematicians are starting to realize that "any time you cubulate something, you're going to reveal all kinds of structural secrets."

The End of an Era

Agol's proof of Wise's conjecture was a four-for-one deal: it proved not just the virtual

Haken conjecture, but the other three of Thurston's 23 questions that were still unresolved. In the years leading up to his proof, Agol and other mathematicians had shown that all three of these questions — the virtual fibering conjecture and two more technical questions about hyperbolic three-manifolds — were also consequences of Wise's conjecture.

In the case of the virtual fibering conjecture, recall that the goal was to show that every compact hyperbolic three-manifold has a finite cover that fibers over the circle, meaning that it is built by gluing the opposite ends of a thickened surface. We know from the virtual Haken theorem that the manifold has a finite cover that is Haken — that is, the manifold's cover has an embedded, incompressible surface. If you cut the Haken manifold open along that surface, you get something that looks like a thickened surface at its ends but has who-knows-what topological features in its "guts."

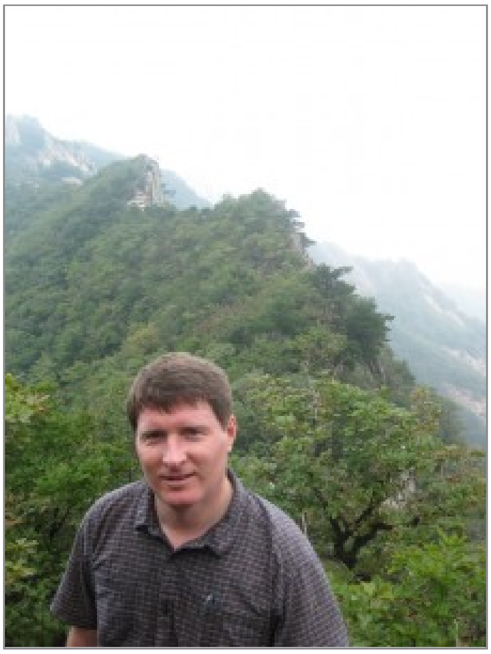

Sang-Hyun Kim

Ian Agol on a recent trip to Daejeon, South Korea.

In 2008, in what Calegari calls "an astonishing breakthrough,"Agol showed that hyperbolic three-manifolds that satisfy a certain technical condition are guaranteed to be virtually fibered. The following year, Wise built on this finding to show that all Haken manifolds are virtually fibered; that is, there is a way to unroll a Haken manifold to produce a finite cover that opens up the complicated topology of the guts, resulting in a simple fibered manifold. Thus, if a manifold is virtually Haken, then it also must be virtually fibered.

"I think everyone had believed that the virtual Haken conjecture would turn out to be true, but the virtual fibering conjecture had seemed orders of magnitude farther out of reach," Calegari said. "To me, the fact that the virtual fibering conjecture follows from the virtual Haken conjecture is one of the most shocking aspects of the story."

With the proof of the virtual fibering conjecture, "you're tempted to think that this means three-manifolds are really simple, because manifolds that fiber over the circle are simple," Minsky said. "But I think it teaches us that manifolds that fiber over the circle are not simple after all — they're more subtle than we expected."

At the same time, the virtual fibering theorem does mean that there is a simple and informative recipe for generating all compact hyperbolic three-manifolds: start with a thickened surface, glue its inner and outer boundary surfaces to each other with your choice of twist, and then fold that manifold over itself a finite number of times.

"If you were to ask me for a hyperbolic three-manifold, I'd ask what kind you want — what kind of fibration and which finite covering?" Calegari said. "We know now that we're not missing out on any three-manifolds by doing this."

While it will take some time for mathematicians to check Agol's work thoroughly, many are optimistic that it will stand up to scrutiny.

"Ian Agol's not a sloppy guy," Minsky said.

Now that the final questions on Thurston's list have presumably been laid to rest, researchers are already starting to ask what the field of three-manifold topology will look like in this brave new post-Thurston world.

Mathematicians agree that they will have plenty to do in figuring out what insights Wise's cube complexes have to offer for other shapes that can be cubulated. When it comes to three-manifolds themselves, mathematicians have reached the end of an era, Agol said, but also the beginning of a new one.

"Most fields of mathematics don't have an all-encompassing vision to guide the field over twenty or thirty years, the way we've had," he said. Now, he suggests, three-manifold topology and geometry may become more like those other fields, in which mathematicians "grope around" and manage to make progress even without the benefit of a big conjectural picture of what is going on.

"New generations of mathematicians will figure out what the next important questions are,"Agol said.

This article was reprinted on ScientificAmerican.com.