


TOPOLOGY

Surfaces Beyond Imagination Are Discovered After Decades-Long Search

By LEILA SLOMAN

June 2, 2022

Using ideas borrowed from graph theory, two mathematicians have shown that extremely complex surfaces are easy to traverse.



One way to describe both surfaces and graphs is to ask how long it would take for someone to randomly walk across the entire object.

—

Myriam Wares for Quanta Magazine

Last July, two mathematicians from Durham University, [Will Hide](#) and [Michael Magee](#), confirmed the existence of a much sought-after sequence of surfaces: each more complicated than the last, ultimately becoming so intricately connected with themselves that they nearly reach the limits of what's possible.

At first, it wasn't obvious these surfaces existed at all. But since the question of their existence first arose in the 1980s, mathematicians have come to realize that these surfaces may actually be commonplace, even if they're exceedingly difficult to pinpoint — a perfect example of how mathematics can subvert human intuition. The new work is a step forward in a quest to move beyond intuition to understand the myriad ways surfaces can manifest.


“It's a brilliant piece of mathematics,” said Peter Sarnak, a mathematician at the Institute for Advanced Study in Princeton, New Jersey.

Surfaces include all kinds of two-dimensional objects — the outer shell of a sphere, doughnut or cylinder; a Möbius strip. They're essential to mathematics and physics. But although mathematicians' relationship with surfaces goes back centuries, they don't know these objects well at all.

Simple surfaces are not the problem. Simple in this case means the surface has a small number of holes, or a low “genus.” A sphere, for example, has no holes, and thus has a genus of zero; a doughnut has one.

But when the genus is high, intuition fails us. When Alex Wright, a mathematician at the University of Michigan, tries to visualize a high-genus surface, he ends up with holes arranged in a tidy row. “If you wanted me to be a bit more creative, I could wrap it around into a circle with many holes. And I'd be hard-pressed to come up with any mental picture fundamentally different from those,” he said. But in high-genus surfaces, holes overlap with one another in complicated ways that make them difficult to grasp. A simple approximation is “as far from representative as it could ever be, in every sense,” Wright said.

This struggle is to be expected, said Laura Monk, a mathematician at the University of Bristol. “You can often make things which are bad. However, making things which are good, which are like the thing we expect typically to be true, is a bit more difficult,” she said.



Michael Magee, a mathematician at Durham University, proved a conjecture about complex surfaces that dates back to 1984.

—

Victoria Greener


That means mathematicians wishing to truly understand the space of surfaces need to find ways to uncover objects that they don't even know exist.

In their July paper, Hide and Magee did just that, confirming the existence of surfaces that mathematicians had wondered about for decades. The conjecture they proved, and the history surrounding it, took inspiration from a different field of math altogether: graph theory.

The Maximum Possible

Graphs, to mathematicians, are networks consisting of points, or nodes, connected by lines, or edges. As early as 1967, mathematicians such as Andrey Kolmogorov were studying networks that impose a cost on connecting two nodes. That led to an example of what would later be called an expander graph: one that keeps the number of edges low, while maintaining high connectivity among nodes.

Expander graphs have since become crucial tools in math and computer science, including in practical areas like cryptography. Like a well-designed highway system, these graphs make it easy to travel from one node to another without covering the entire graph with edges. Mathematicians like to limit the number of edges by stipulating that each node can only have, say, three edges emanating from it — just as you might not want more than a few highways crisscrossing your town.



If a computer chooses at random where each node's three edges lead to, you'll find that — especially when the graph is very large — most of these random graphs are excellent expanders. But though the universe is filthy with expander graphs, human beings have failed time and again to produce them by hand.

“If you want to build one, you shouldn't draw them yourself,” said [Shai Evra](#), a mathematician at the Hebrew University of Jerusalem. “Our imagination does not grasp what is an expander.”

The idea of expansion, or connectedness, can be measured in a number of ways. One is to cut a graph into two large pieces by snipping edges one by one. If your graph consists of two clumps of nodes, with the clumps connected by a single edge, you only need to cut a single edge to split it into two. The more connected the graph, the more edges you'll have to slice through.

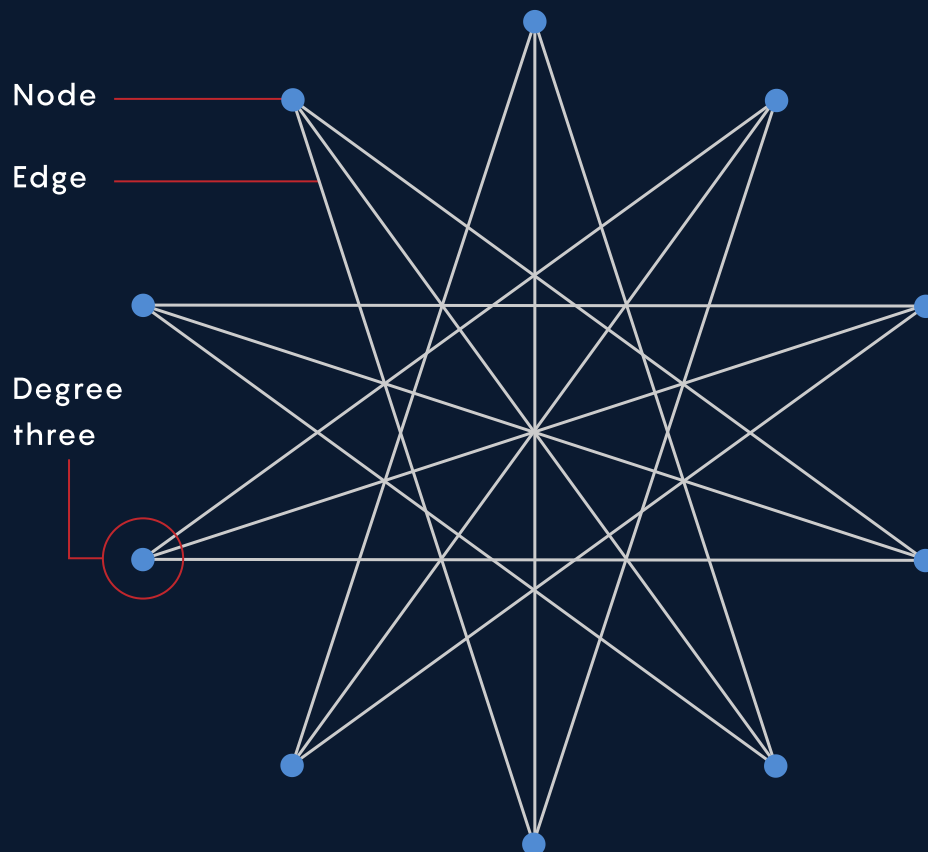
Another way to get at connectivity is to wander around the graph from node to node, at each step choosing an edge to walk along at random. How long will it take to visit all the graph's neighborhoods? In the example with the two clumps, you'll be confined to one of the bubbles unless you happen to step across the lone connection to the other half. But if there are many ways for you to travel between different areas of the graph, you'll tour the whole thing in a short amount of time.

Connected or Not?

Graphs consist of points, or nodes, connected by lines, or edges. If a graph is well connected, you'll need to cut many edges to break it in half.

EXPANDER GRAPHS


Have relatively few edges — here, just three per node — but a high degree of connectivity.



POORLY CONNECTED GRAPHS

Have few paths from one node to another.





Merrill Sherman/Quanta Magazine

These measures of connectedness can be quantified by a number called the spectral gap. The spectral gap is zero when the graph is completely disconnected, for instance if it's made up of two groups of nodes that aren't attached to each other at all. As a graph gets more connected, its spectral gap will tend to get bigger.

But the spectral gap can only go so high. Indeed, the two defining features of expander graphs — few edges and high connectivity — are seemingly at odds with one another. But in 1988, [Gregory Margulis](#) and, independently, Sarnak and two co-authors [described](#) “optimal expanders” — graphs whose spectral gap is as high as the theoretical maximum. “It’s shocking that they exist,” Sarnak said.

Later, mathematicians would prove that most large graphs are close to this maximum. But the work with optimal expanders and random graphs was not a matter of just finding the right places to put edges. It required the use of strange and sophisticated techniques borrowed from number theory and probability.

Applying those results to surfaces would require a stranger approach still.

Surface Tension

In the 1970s, at the same time that graph theorists were mapping out expander graphs, a Swiss geometer named Peter Buser was playing with similar ideas in surfaces with lots of holes. Could they be well connected in the same way expander graphs were?

As with graphs, you can estimate a surface's connectedness by attempting to slice it in two, or by wandering aimlessly around the surface and seeing how long it takes you to visit the whole thing. “Everything is undeniably more sophisticated, but in every way the intuition goes through,” said Wright.



The project given to William Hide, a doctoral student at Durham, ended up being far more substantial than expected.

Victoria Greener

In 1969, Jeff Cheeger proved that the spectral gap on surfaces is related to how easy it is to cut the surface in half. Buser later strengthened Cheeger's result, making the spectral gap a handy barometer of connectivity. Buser began to wonder what happens to the spectral gap as surfaces get more and more complicated — whether surfaces with many holes can still be well connected.

At first, he guessed that the spectral gap would have to shrink to near zero for high-genus “hyperbolic” surfaces, an important class of surfaces that are curved in a particular way. But as he learned about the emerging examples of expander graphs, he changed his thinking. “Gradually, I learned that this is not true,” said Buser, who is now a professor emeritus at the Swiss Federal Institute of Technology Lausanne. “I just had to learn, no, no, your intuition is wrong.”

In 1974, Buser's doctoral adviser had found the ceiling for a spectral gap on a surface: just above $1/4$, with the exact number depending on how many holes the surface has. (The more holes, the lower the number.) Buser wanted to know more about this limitation. In particular, were there any real surfaces that reached the cutoff? Was there a geometric equivalent to optimal expander graphs?

In 1984, Buser showed that some hyperbolic surfaces with a large number of holes have a spectral gap of at least $3/16$. In that paper, he noted: “It is conjectured that $3/16$ can be replaced by $1/4$.”

Nearly four decades on, mathematicians have finally proved that statement true.

Probably a Good Idea

Hide and Magee's project didn't start out as an attempt to reach $1/4$ at all. When Hide began as Magee's doctoral student in the autumn of 2020, Magee suggested what he thought would be a manageable project for a new doctoral student: Hide would try to reproduce one of Magee's recent results, but applied to a new situation.

In the spring of 2020, Magee, Frédéric Naud of Sorbonne University and Doron Puder of Tel Aviv University had found that many “compact” surfaces have a relative spectral gap approaching $3/16$. Compact surfaces, which include spheres and doughnuts, have two important properties. They're a finite size — unlike, say, the infinite plane — and they have no edges. Buser's studies of the 1970s and '80s, as well as his musings about the spectral gap, had all been focused on compact surfaces. Magee and his co-authors dreamed of moving their result from $3/16$ all the way up to the conjectured $1/4$. But the path there seemed almost as obscure as it was in 1984.

Magee asked Hide to think about a type of surface much like a compact surface, but with long, slender spears poking out of it. These are called cusps, and they become so thin at infinity that their surface area is still finite, even though they're infinitely long. That makes these surfaces “finite-area non-compact.”

Hide and Magee thought they could update ideas from Magee's earlier work, which involved generating surfaces randomly. These kinds of random strategies have been used to great effect in mathematics, particularly in graph theory. They allow mathematicians to get a firm understanding of the typical properties of an object. “Probabilities allow you to throw away bad examples,” said Monk. If you compute the probability that a marble is round and get 99%, you've just gotten a lot of information

about what marbles are like, without getting distracted by the reality that some marbles are chipped or deformed.

But the two mathematicians struggled to make the strategy work. They couldn't use random strategies to say anything substantial about the objects they were interested in. In the spring of 2021, they gave up — at least, for a few months.

In June, Naud brought up something Magee thought he and Hide could use: In 1988, Buser and two co-authors had come up with a way to manipulate a finite-area surface into something compact.

Suddenly, Hide's abandoned doctoral project took on a new meaning. "If you could prove not just $3/16$, but also $1/4$, this would also resolve [the compact] problem," said Magee.

With renewed vigor, and armed with ideas that had worked in the infinite-area case, Hide and Magee returned to their original problem with the more ambitious goal of finding a sequence of finite-area non-compact surfaces whose spectral gap approached not just $3/16$, but $1/4$.

This time, they succeeded, posting their paper online on July 12.

"It was really surprising that this thing I gave my Ph.D. student I thought was some easy thing, it turned out to be really important," said Magee. "More important than I ever thought it would be."

All the Surfaces

Though Hide and Magee resolved a long-standing mystery, there are many directions for mathematicians to explore from here. One possibility is to try to prove there are high-genus surfaces whose spectral gap is actually the maximal value, $1/4$, rather than arbitrarily close to it.

In addition, Hide and Magee's technique only works for a very specific type of surface, meaning any probabilities they calculate will only tell you about a tiny fraction of all the surfaces out there.

But there is a model out there, the Weil-Petersson model, which applies to the entire space of hyperbolic surfaces — and which researchers have already studied in this context, most notably Maryam Mirzakhani. Her work on this model, said Monk, "was a complete revolution creating a whole field." The Weil-Petersson approach, however, is still stuck at $3/16$.

Whatever the future holds, the Hide-Magee result is worth celebrating on its own, mathematicians say. "We're asking very precise results for objects that are generally elusive. It's very difficult," said Hugo Parlier, a mathematician at the University of Luxembourg. "I think we should be more amazed by the fact that we're able to say anything at all about the [spectral gap] that's so precise."