Kadison–Singer Problem Solved

By Dana Mackenzie

What is the best kind of mathematical problem? Does it resemble a distant mountaintop, beckoning people from far and wide to attack the summit, simply "because it's there"? Or is it like a vast subterranean river, connecting different realms of mathematics, mysteriously disappearing below the surface and reappearing where you least expect it?

Proponents of either kind of problem have a good case. But this year, it was the turn of one of the great subterranean rivers of mathematics to emerge into the limelight. In June, a three-member team of mathematicians and computer scientists posted a proof online for the Kadison–Singer conjecture (sometimes called, more accurately, the Kadison–Singer problem). It's a problem whose source lies in the forbidding uplands of C*-algebras and quantum physics. But through the 1990s and 2000s, it kept surfacing in other places. It meandered through operator theory, complex analysis, graph theory, and signal processing, before finding a formulation as a seemingly simple problem in finite-dimensional geometry.

As Dan Spielman of Yale University explains it, he got interested in the Kadison–Singer problem partly by accident. In 2009, he published a paper with two former students, Nikhil Srivastava and Joshua Batson, on sparsification of graphs. Graphs, in this context, are networks, most likely very large and somewhat random. A good example is the Internet or Facebook. To understand networks, graph theorists ask questions like: What are the cliques? If the edges have distances or weights on them, what is the shortest path that visits every node (the traveling salesman problem)? These are very difficult questions to answer for a typical dense network, with a large number N of nodes and an even larger number (proportional to N^2) of links. It would be very useful for certain applications to replace a dense network with a simpler approximation—a network whose total number of links is proportional to N but that would nevertheless have a similar overall structure.

A widely used tool for measuring the similarity of graphs is the spectrum of the Laplacian operator of the graph. This is a discrete analogue of the differential Laplacian operator that mathematicians and physicists have used for two centuries to describe objects in the plane, in space, and in higher dimensions. The classical problem of "hearing the shape of a drum" asks what the eigenvalues (or spectrum) of the Laplacian reveal about the shape of a space. Analogously, the goal of sparsification is to create a network that "sounds" like another one but is much simpler.

Sparse approximations of the "complete graph," a network of N vertices in which every vertex is connected to every other, had been known for some time. These approximations have proved incredibly useful in computer science and engineering; for example, they have been used to construct error-correcting codes, hash functions for cryptography, and pseudo-random number generators. Spielman, Srivastava, and Batson had generalized this result to show the existence of approximations for any graph. Their usefulness, however, was limited by the fact that the researchers could not control the weights on the edges.

When Spielman told him about the work, Gil Kalai of the Hebrew University of Jerusalem remarked that their theorem sounded a lot like one version of the Kadison–Singer problem. In fact, the Kadison–Singer conjecture, if true, would give Spielman's group the control over the edge weights that had been missing.

A couple months later, Adam Marcus arrived at the Yale mathematics department; a newly minted PhD, he was looking for a new project to work on. Spielman suggested the Kadison–Singer problem, which, he says, got Marcus "very excited." The two of them started collaborating with Srivastava, who now works at Microsoft Research in India. Little did they suspect that it would take five years of hard work to get an answer. "But they were rewarding years," Spielman says. "We kept feeling as if we were making progress. It was like a computer video game, giving you just enough information to keep you working on it."

The original version of the Kadison–Singer problem looks about as unrelated to graph theory as a problem could possibly be. What Richard Kadison and Isadore Singer asked in 1959 is whether "pure states" on abelian von Neumann algebras could be extended uniquely to pure states on non-abelian algebras.

The question almost defies translation into simple English, but let's give it a shot. We can think of the elements of a von Neumann algebra as infinitely large matrices. In quantum physics, they would be observable quantities, such as position, momentum, or spin. A state is a function of these matrices; in physics, we can think of it as the probability of measuring a certain value of the observable in question. When the algebra is "abelian" (i.e., the product of two matrices A and B is the same in either order), the matrices correspond to observables that can be measured simultaneously. But by Heisenberg's uncertainty principle, some observables, such as position and velocity, cannot be measured simultaneously. These correspond to matrices that do not commute. Kadison and Singer asked whether a certain kind of state defined on a commuting set of matrices can be extended in a unique way to a larger non-commuting set. To put it in physical terms: Can a series of measurements of one observable uniquely constrain the other observables that can't be measured at the same time?

Of course, technical details need to be worked out. "You're talking about non-normal pure states, non-physical states, which only exist by the axiom of choice," says Nik Weaver, a functional analyst at Washington University in St. Louis. (The axiom of choice is a well-known assumption that often leads to very non-intuitive results, but is also indispensable for reasoning about infinite-dimensional spaces.) For so-called "normal" states, which might actually be observed in a physics experiment, the answer to the Kadison–Singer question is yes, as Kadison and Singer proved. They suspected, however, that the answer for the non-normal states was no. With admirable foresight, they did not actually state this as a conjecture. As it turned out, their guess would have been wrong.

In its original formulation, the problem was of great interest to specialists in operator algebras, but rather inaccessible to anyone else. Beginning around 1979, however, the underground river began to well up in different places and under different names: the Paving conjecture, the Bourgain–Tzafriri conjecture, the Feichtinger conjecture. The mathematicians who posed these problems were not always aware of each other, and it wasn't until 2006 that Peter Casazza of the University of Missouri (with three other mathematicians) put them all together. "We didn't start out looking for equivalent versions of Kadison–Singer," Casazza says. "We were actually trying to solve the problem, and kept moving to different areas of research, hoping that one of them had deep enough results to handle the problem. It was a fluke that each time we entered a new area of research, this was equivalent to their most famous unsolved problem."

Almost imperceptibly, however, the river was approaching its destination—the conjecture was becoming simpler to state, if not to prove. A major step forward was the discovery of finite-dimensional versions of the problem. One of Casazza's favorite versions involves finite sets of vectors called "frames."

Linear algebra students learn in college that an "orthonormal frame" of vectors $(\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_N)$ can be used to reconstruct any vector $\mathbf{x}$ from an N-dimensional space, as follows:

$$\mathbf{x} = (\mathbf{x} \cdot \mathbf{v}_1)\mathbf{v}_1 + (\mathbf{x} \cdot \mathbf{v}_2)\mathbf{v}_2 + \ldots (\mathbf{x} \cdot \mathbf{v}_N)\mathbf{v}_N.$$

In fact, many other collections of vectors—with far more than N elements—also enable the reconstruction of $\mathbf{x}$ via the same formula: in some cases exactly, in others approximately.

If the vector $\mathbf{x}$ is thought of as a signal, frames make it possible to reconstruct $\mathbf{x}$ from a redundant set of measurements. The question that arises then is: What if data are lost in transmission? Can $\mathbf{x}$ be recovered from subsets of the frame? The Kadison–Singer conjecture says yes. Any frame can be partitioned into smaller frames, with uniform bounds both on the number of smaller frames and on the loss of fidelity.

In 2013, Spielman, Marcus, and Srivastava finally proved a closely related conjecture of Weaver, which involves partitioning a set of vectors into just two "balanced" subsets. The balancing condition amounts to saying that the largest root of a certain polynomial will be less than 0.9 (or some fixed constant less than 1). In fact, Spielman's team proved that the largest root of the *average* polynomial over all possible partitions is less than 0.9.

"It's like saying that the average family has 2.5 children," Spielman says. "If the average family has 2.5 children, you know that some families have two or fewer." Similarly, one might think that if all the roots of the average polynomial for all partitions are less than 0.9, then there must be at least one specific polynomial with that property.

In fact, that was what Spielman's group thought at first. But, he says, "It was a mistake." There are no general relationships between the roots of an average of a set of polynomials and the roots of the individual polynomials. But the group discovered one case in which they are related—namely, when the roots of the polynomials "interlace." The researchers had to build up the theory of interlacing polynomials essentially from scratch, but ultimately it allowed them, as Spielman says, "to use the average to reason about individuals."

Their highly innovative result applies immediately to graph theory. For example, as Nick Harvey of the University of British Columbia points out, one promising approach to the traveling salesman problem is to find "optimal thin trees" for a given graph. These are sparse subgraphs with no loops. The Kadison–Singer theorem proves that optimal trees exist that *look* like good approximations from the perspective of the Laplacian spectrum. If the "optimal spectrally thin trees" could be converted to optimal thin trees (admittedly a big if), it would be a major breakthrough on the traveling salesman problem.

The result could also have applications to signal processing. If your cell phone works better 20 years from now, you might have Kadison and Singer (and Spielman, Marcus, and Srivastava) to thank. But one major hurdle needs to be overcome between now and then. The proof is non-constructive, which means that it does not give a recipe for computing balanced partitions, or any of the other equivalent objects whose existence it promises. "If you want to turn the proof into an algorithm, it would take exponential time," Spielman says. In other words, as a practical search method the proof is not much better than trial and error. In future years, it will be a major project for computer scientists to translate the theoretical existence theorem into a practical tool for engineers. Says Casazza, "It's the engineers, in the end, who will decide whether this is really valuable."

Dana Mackenzie writes from Santa Cruz, California.