N news

Science | DOI:10.1145/3386366

Don Monroe

A Proof from 'The Book'

A decades-old conjecture about computational complexity is confirmed in just a few pages.

IN 1637, THE French mathematician Pierre de Fermat scribbled in a book that he had a proof for a theorem, which the margin was too small to contain. It seems likely that he was mistaken about his famous "last theorem," since no such proof was found for 358 years, and even then it required more than 100 pages and used mathematics that didn't exist in his time.

By contrast, the Boolean "sensitivity conjecture" is relatively recent, but after nearly 30 years and repeated failures it seemed likely that any proof would also be a long and difficult slog. In July, 2019, however, mathematician Hao Huang of Emory University in Atlanta posted a short paper that completed its proof in a couple of pages in a way that experts found very convincing. Indeed, they immediately described it as "from the book," a tome imagined by the great 20th-century mathematician Paul Erdös in which God records the shortest and most illuminating proof of each theorem.

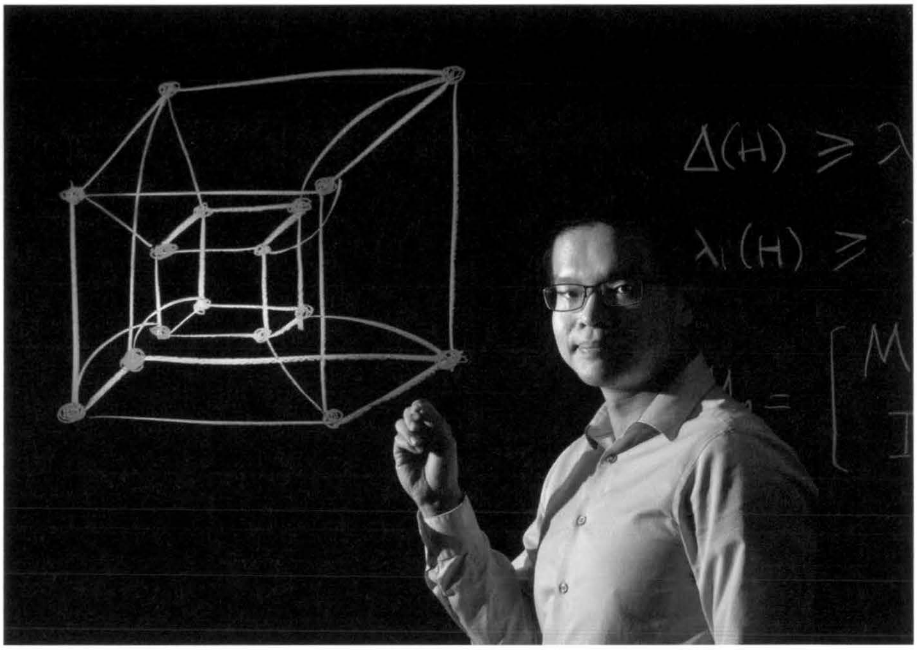

Emory University mathematician Hao Huang, who developed the proof of the sensitivity conjecture.

The new proof closes a nagging loophole by confirming that "sensitivity" of any Boolean function is closely related to other measures of its computational complexity. "In the past, computer scientists studied a lot of different complexity measures. All but one of them were known to be polynomially related," Huang said, meaning that their possible values are constrained by powers of the others. "My work is basically to put the sensitivity—the last exception—into this category."

Algorithmic complexity theory aims to provide strict upper and lower bounds on the difficulty of calculations, especially as the problems get larger. Its best-known open question, P≠NP, concerns the existence of problems whose solution can be verified with computational resources that grow as a polynomial function of the problem size but which require much more effort to solve in the first place. The new result is "not going to spring open P≠NP," cautioned Kenneth Regan of the University at Buffalo (part of the State University of New York system). Nonetheless, he noted that the sensitivity conjecture is closely related to "the tools that people have been using to try to get a handle on P≠NP."

Measures of Complexity

Boolean functions are at the heart of digital computation, producing a one-bit (zero or one) output, based on the values of a string of input bits. Sensitivity is one measure of a Boolean function's complexity: Considering all possible input strings, the sensitivity is the largest number of input bits whose individual flipping (from zero to one or vice versa) changes the output.

Sometimes flipping any input bit is enough. For example, the parity function, which reports whether the input has an even or odd number of ones, changes value for any input bit, for any input string. Thus its sensitivity has the maximum possible value, equal to the number of input bits. The logical "and" of all the bits also has maximum sensitivity, because, for an all-ones input, then changing any bit to zero flips the output.

Other functions are less sensitive, so that for all input strings there are some bits that cannot on their own change the answer. For example, the OR-of-AND function asks, for a bunch of non-overlapping blocks of input bits, whether any of them has all ones for inputs. Changing the answer from yes can only be done by flipping a bit in a unique all-ones block, while changing the answer from no can only be done by flipping the last remaining bit that is not zero in one of the blocks. The sensitivity is the larger of the number of bits per block or the number of blocks.

Computer scientists have explored several other measures of computational complexity. The "block sensitivity," for example, quantifies how many blocks containing multiple input bits change the output when they are simultaneously flipped. Another measure is the function's "degree," which is the highest total exponent in the polynomial that reproduces the output when the inputs are zero or one.

Other measures include decision-tree depth (the minimum number of yes-or-no questions needed to guarantee knowing the output), as well as its quantum and random variants, and certificate complexity (the number of input bits needed to guarantee knowing the output). Researchers proved long ago that all these other complexity measures are closely related. Specifically, upper or lower bounds on their values, for any Boolean function, can be expressed as polynomials of the others, which is useful for proofs. "If they are polynomially related, they are roughly of the same order of magnitude," Huang said, "so instead of looking at the more difficult ones, you can look at a simpler one."

ACM News

The Shuttering of Corporate Datacenters

Corporate datacenters are being decommissioned rapidly. In a blog post, market research firm Gartner forecast that 80% of enterprises will have shut down their traditional datacenters by 2025.

Many companies are now migrating their in-house datacenters to the cloud. Oracle predicts 80% of enterprise workloads will move to the cloud by 2025.

According to Richard Villars, vice president for datacenter and cloud at market research firm IDC, the cloud is only one factor contributing to datacenter decommissioning. "Virtualization and converged infrastructure allowed many corporations to get a lot more capacity in their datacenters," Villars said. "With these technologies, you could now do the same amount of work on 40 or 50% of the footprint."

Bill Vasquez, senior vice president of Strategy & Business Development at ITRenew, which specializes in onsite datacenter decommissioning and data erasure services, agrees decommissioning is growing at a rapid pace. One of the main growth drivers is the sheer volume of hardware deployed, he says.

Companies retain servers for four years on average, and the number of servers in a datacenter can run as high as 80,000. IDC reports enterprise computing is at near-historic highs, with next-generation workloads and advanced server innovation driving server demand at the end of 2019 to one of the highest levels in 16 years, according to a recent IDC Quarterly Server Tracker. It is the deployment of these new servers that is spurring the decommissioning of the older equipment they are replacing.

"Between the growth in data usage and storage, and the emergence of new technologies that require more and more computing power, like artificial intelligence, machine learning, augmented & virtual reality, and the Internet of Things," Vasquez says. "Decommissioning shows no sign of slowing down."

Considering all the data that resides on the servers from decommissioned datacenters, "The best way to verify data has been destroyed is to wipe it with 100% sector-verified erasure, and electronically capture the serial number of both the host unit and the media itself with a solution like Teraware," Vasquez says, pointing to ITRenew's data-wiping software. Wiped drives can be reconciled against an asset inventory system for further verification, he adds, and enterprises in industries with higher security requirements may require the units to be shredded after they have been wiped.

"How we actually destroy the media that has the information on it is where the rubber meets the road from a certitude perspective," says Bob Johnson, CEO of the National Association of Information Destruction (NAID). "Some data centers' internal IT staff may do their own wiping and disassembly before disposal, but in a large percentage of cases, they are turning over their equipment and relying on a third party to perform the data destruction, as well as the equipment removal and recycling."

NAID verifies secure data destruction companies' services compliance with data protection laws through audits by accredited security professionals, fulfilling customers' regulatory due diligence obligations.

The data owner is always responsible for the protection of its data, as well as for regulatory compliance. "They are not able to contract that away," Johnson says.

Vasquez agrees the owner of the hardware is ultimately responsible for the data and its destruction, which is why it is of paramount importance for them to complete thorough due diligence when deciding on a potential data destruction partner. "Only partners who provide the most stringent security solutions should be trusted with this work," he says.

—John Delaney is a freelance writer based in New York City, NY, USA.

Coloring Hypercubes

Some important conclusions have been mathematically expressed only in terms of sensitivity, however, and until now there had been no proof that they were also members of this club, although it was widely thought to be. Indeed, in 1992, Noam Nisan of the Hebrew University of Jerusalem and Mario Szegedy, then at AT&T Bell Laboratories, explicitly conjectured that the block sensitivity of a function could not exceed some fixed power of its sensitivity.

A critical strategy for proving this "sensitivity conjecture" was provided the same year by Craig Gotsman and Nathan Linial, both then at the Hebrew University of Jerusalem, who connected the sensitivity with the graph-theoretical properties of corners of a hypercube.

The coordinates of any corner of an n-dimensional hypercube can be written as an n-bit string of zeroes and ones. A Boolean function then corresponds to coloring the corners, say red when the function is one and white when it is zero.

If exactly half of the corners are red (and half white), they can be arranged so that no corner has a like-colored neighbor, for example by coloring them according the parity function of their coordinates. If even one additional corner is colored red, however, it turns out that at least one of the red corners must have many red neighbors. The question is how many? The largest number, among all red corners, is called the degree of the framework (which is not the same as the degree of the function).

In this picture, the sensitivity of a function is the maximum number of white corners (opposite output) sharing an edge (one input-bit flip) with any red point. The remainder of the n edges are connected to red points, so the sensitivity is closely connected to the subgraph's degree.

The half-page proof by Gotsman and Linial showed roughly that a bound on a subgraph's degree, as a function of n, is equivalent to a bound on the sensitivity of a Boolean function, as a function of the degree of its polynomial. Thus, a theorem about one becomes a theorem about the other, opening the door to a proof of the sensitivity conjecture.

The proof "provides a useful addition to the toolbox of mathematics and computer science that hopefully will see more application in the future."

Changing Signs

In spite of this clear roadmap, and decades of attempts, this promise was only realized with Huang's proof. What Huang showed was that if even one more than half of the hypercube's corners are red, at least one of them will have at least $\sqrt{n}$ red neighbors, precisely what Gotsman and Linial suspected. This implies that the degree of the function is no greater than the square of its sensitivity. Together with a previous theorem that the block sensitivity is no greater than the square of the degree, this means that the block sensitivity does not exceed the fourth power of the sensitivity. This confirms the conjecture and thus connects sensitivity to all the other complexity measures.

The proof is based on something called the adjacency matrix, whose 2^n rows and 2^n columns correspond to the corners of the hypercube, and whose elements are zero unless the corners are same-colored neighbors. The critical trick is that the non-zero elements for neighbors are usually assigned a value of +1, but Huang assigned some of them a value of −1. Carefully choosing the negative values allowed him to use known matrix theorems to show that at least one row of the matrix must have at least $\sqrt{n}$ entries.

Aaronson wrote on his blog, "How could such an elementary 1.5-page argument have been overlooked for 30 years? I don't have a compelling answer to that, besides noting that 'short' and 'elementary' often have little to do with 'obvious.' Once you start looking at … this matrix …, the pieces snap together in precisely the right way—but how would you know to look at that?"

"For making a calculation it's straightforward to go through every step," Huang noted. In contrast, "a proof is easy to verify, but it's difficult to come up with a new proof."

Since encountering this problem in 2012, Huang wrote in a comment on Scott Aaronson's blog, "I revisited this conjecture every time I learned a new tool—without any success, though."

More recently, he added, "I had been looking at other kinds of problems at the same time, and I used this adjacency matrix a lot," he said. "I realized that it can be also applied to this particular sensitivity conjecture, and that's how I came out with the proof."

Regan, who had previously explored similar ideas, said the cancellations allowed by making some matrix elements negative is critical, and are reminiscent of the interference that quantum algorithms exploited.

Huang said the proof "provides a useful addition to the toolbox of mathematics and computer science that hopefully will see more application in the future," although he noted that in math this process often takes many years. He also hoped that the work would be an inspiration for graduate students to attack long-standing unsolved problems. ■

Further Reading

Huang, H.,
Induced subgraphs of hypercubes and a proof of the Sensitivity Conjecture, July 1, 2019 https://arxiv.org/abs/1907.00847

Klarreich, E.,
Decades-Old Computer Science Conjecture Solved in Two Pages, *Quanta*, July 25, 2019, http://bit.ly/36ipHTo

Gotsman, C., and Linial, N.,
The equivalence of two problems on the cube, *Journal of Combinatorial Theory, Series 1*, Volume 61 Issue 1, September 1992, pp. 142-146.

Blog Posts:

Sensitivity Conjecture Resolved, Scott Aaronson at *Shtetl-Optimized*, July 2, 2019.

Tools and Sensitivity, Ken Regan at *Gödel's Lost Letter and P=NP*, July 12, 2019.

Amazing: Hao Huang Proved the Sensitivity Conjecture! Gil Kalai at *Combinatorics and More*, July 2, 2019

Don Monroe is a science and technology writer based in Boston, MA, USA.

© 2020 ACM 0001-0782/20/5 $15.00