




A Satisfying Result

Formulating a decades-old geometric conjecture as a satisfiability problem opened the door to its final resolution.

A TEAM OF RESEARCHERS has completed the understanding of the Keller Conjecture, first proposed in 1930, about the packing of squares, cubes, and their higher-dimensional analogues. The conjecture states that any tiling of identical hypercubes that fills space must contain a pair of neighbors that share an entire face.

You might convince yourself it is true for two or three dimensions by toying with squares or blocks. Mathematicians later established the status of the conjecture for all dimensions except seven. The new result, which shared the best paper award at the 2020 International Joint Conference on Automated Reasoning (IJCAR 2020), fills that gap. To do this, the researchers mapped more than 10,324 ways seven-dimensional hypercubes can avoid sharing any six-dimensional face onto a satisfiability problem, which asks whether some Boolean formula can ever be made true.

After exploiting the symmetries of the problem to reduce it to a mere billion or so distinct configurations, the team used extensive computer resources to systematically eliminate them all, implying there must always be a shared face. The resulting proof is enormous, about 200 gigabytes. Obviously, it could not possibly be checked by humans,



A graph representing the two-dimensional case of the Keller Conjecture.

but it produces a certificate that can be checked by other computer programs, and so is almost certainly correct.

"It is satisfying because it's finishing off an old problem with a venerable history, and math advances like that," said Jeffrey Lagarias of the University of Michigan, whose work nearly 30 years ago disproved the most general framing of the conjecture. The proof also represents another success for satisfiability-solver algorithms, which are solving problems in many areas of computer science and mathematics.

A Long History

Mathematician Hermann Minkowski proposed the face-sharing rule for tilings in which every hypercube sits on a periodic lattice. His conjecture was proven in all dimensions in the 1940s. Eduard Ott-Heinrich Keller's generalization to non-lattice tilings (which includes periodic arrangements of clusters of hypercubes, rather than individual hypercubes) was also proven up to dimension six; higher-dimensional non-lattice tilings, however, remained in play until 1992, when Lagarias and Peter Shor (both then at AT&T Bell Labs) found a non-face-sharing counterexample in 10 dimensions. In a five-page paper, they described a tiling of 10-dimensional hypercubes, none of which share a complete nine-dimensional face.

That result showed that Keller's original conjecture, which covered all dimensions, was false. Moreover, a non-face-sharing tiling in a particular dimension can be used as a layer in the next-higher dimension, with a simple slide between neighboring layers preventing face-sharing. Thus, once the conjecture was violated in 10 dimensions, it was ruled out for all higher dimensions.

Ten years later, John Mackey, then at Harvard University, found a counterexample for eight dimensions (and thus for nine). Until now, however, mathematicians did not know whether eight was the lowest number of dimensions with a counterexample, or if seven was.

The computer-aided solution built on several intermediate advances. "The interesting question is, how is it that something that appears to be an infinite problem can be reduced to something that can be done by com-

"It is satisfying because it's finishing off an old problem with a venerable history, and math advances like that."

puter," said Thomas Hales of the University of Pittsburgh. In 1998, Hales used exhaustive computer checking for a proof of a famous 1611 conjecture by Johannes Kepler concerning the densest packing of spheres.


For Keller's conjecture, Hales said, "The first reduction that you can do is bring it down to a periodic arrangement." It turns out that any non-face-sharing example would have to comprise a lattice of clusters that are no larger than 2 in each of the d dimensions, comprising 2^d cubes.

"That's not enough to reduce it to a finite problem, because each cube can slide continuously," noted Lagarias. "You need to do a further reduction so that each cube is just at fixed, discrete positions in space." It turns out that any non-face-sharing pattern must feature shifts that are integer multiples of $1/s$, where s ranges up to $d-1$. "Once there are finitely-many positions, and finitely many cubes, then it's a finite search problem," Hales said. Still, he cautioned, "That doesn't mean that you have a computer that has the capacity to solve the problem."

For the outstanding seven-dimensional question, Polish mathematician Andrzej Kisielewicz and a coworker, in the past few years, eliminated every possible counterexample except those with $s=3$, corresponding to cube shifts by $1/3$ of their size. To eliminate this final possibility, Mackey, now at Carnegie Mellon University in Pittsburgh, teamed up with a new professor there, Marijn Heule, as well as undergraduate Joshua Brakensiek, now at Stanford University, and visitor David Narváez of the Rochester Institute of Technology. Heule had exten-

ACM Member News

CREATING ROBOTS FOR USE IN DISASTERS



Robin Murphy is the Raytheon Professor of Computer Science and Engineering at Texas A&M

University in College Station, TX.

Murphy earned her undergraduate degree in mechanical engineering from the Georgia Institute of Technology (Georgia Tech) in Atlanta, where she also earned her master's and doctoral degrees, both in computer science. Additionally, she was a Rockwell International Doctoral Fellow.

After receiving her Ph.D. in 1992, Murphy became a professor at the Colorado School of Mines, before moving to the University of South Florida. She joined the faculty of Texas A&M in 2008, where she has remained since.

Murphy's research is primarily in artificial intelligence for disaster robotics, specializing in human-robot interactions to help them work better together. Her current research is focused on how robots are being used in the response to COVID-19 throughout the world, which she says is telling us a great deal about how robots are adopted for disasters in general.

"Most disaster responses are over in a span of somewhere between three days to two weeks, but the coronavirus timeline is far longer," Murphy points out. This elongated timeframe provides the opportunity to investigate the maturity of the technology, the influence of trust, and which groups adopt, she adds.

"Instead of just being clinical-care hospitals and the public safety realm, every sector is adopting robots, and using them to cope in some form or fashion," she says.

Murphy enjoys formalizing what she has learned as she moves from empirical observations to quantitative analysis methods. "It's been a fun process to develop," she says.

—John Delaney

sive expertise in large-scale satisfiability proofs, notably the 2016 solution of the “Pythagorean triples problem” which, at 200 terabytes, has the dubious distinction of being the largest proof in history (to date).

Computer Attack

The 1992 proof by Lagarias and Shor for 10 dimensions was a simple description of a 1,024-cube counterexample. The *non*-existence proof for seven dimensions is vastly harder because it must comprehensively eliminate all of the 128-cube clusters that are not counterexamples.

Lagarias and Shor exploited a graph-theory formulation of the tiling problem that had recently been crafted by Hungarian mathematician Sándor Szabó. He and a colleague described nodes known as “Keller dice,” which bear spots of various colors (depending on d and s) and are connected by an edge in the graph if the colors satisfy certain rules. The existence of a non-face-sharing tiling is equivalent, in seven dimensions, to the existence of a 128-node (27) clique of connected nodes in the corresponding graph.

This graph-theoretic existence question is naturally framed as a satisfiability problem. The required properties of the clique are described by a Boolean formula, like “(die 4 has a red dot OR die 4 has a green dot) AND (die 5 has a black dot),” but much longer. A satisfiability solver, or SAT solver, examines all possible colorings of the dots to determine if any combination can make the Boolean formula true; that is, satisfy it.

Even with today’s powerful computers, it is critical to further reduce the problem by exploiting symmetries; for example, the fact that swapping the labels of two dice does not change anything important. “Many of these graph problems, the main problem of solving them is the symmetries,” Heule stressed. Moreover, experience with SAT-solver competitions has shown that “If there is a bug in a solver, it’s typically because they do some symmetry breaking which is not sound.” For this reason, Heule insisted these arguments be included in the formal published description of the proof.

The remaining problem was big, but easily manageable with modern

“This technology, formal verification, is far more reliable than the level of verification done by referees for a journal—much more reliable by order of magnitude.”

supercomputing resources, including the Texas Advanced Computing Center at his former institution, the University of Texas at Austin. “I regularly use 5,000 CPUs,” Heule said, although the Keller work only required about 40.

Although he expects he could access greater resources if needed, Heule also strives to be certain he can finish a job before he starts. “I’ve been investing quite a bit in good runtime estimation,” he said. This is especially important if the formula turns out to be unsatisfiable, which demands examining every possible counterexample. In contrast, Heule said, “For satisfiables, you can get lucky.”

“There are lots of problems that you can translate into satisfiability,” including questions in first-order logic and halting problems like termination of term-rewriting algorithms, Heule said. “More and more fields are seeing the potential of using satisfiability, and that satisfiability is more powerful than any of the dedicated techniques they developed.”

Trust, But Verify

“One really nice thing about SAT solvers now is that they don’t just give you a yes-no answer in the end that you have to trust,” Hales said. “Even if the proof is very large, a proof certificate is produced, and that can be separately checked.” He adds that “This technology, formal verification, is far more reliable than the level of verification done by referees for a journal—much more reliable by orders of magnitude.”

“This resolution of Keller’s conjecture is the absolutely best possible thing,” agreed Lagarias. “You really wanted a computer-aided proof with a proof checker.”

For number theorist Michael Harris of Columbia University, though, “The issue is not so much whether the reasoning is correct as how it works. If you can’t understand the reasoning, it hardly matters whether it’s correct or not, from the point of view of mathematicians.” Proofs “are not generally end-products, they are stages in a historical process,” he stressed. “An impenetrable proof doesn’t provide what mathematicians are looking for,” Harris said, and he is skeptical of “vague claims that at some point computers are going to be better at doing mathematics than human beings, whatever that means.”

Heule emphasizes the ways computers can augment mathematical insight. For one thing, “I think it’s better to have an answer than no answer, although we don’t understand it.” He also described an eminent mathematician, who told Heule “he wastes about 70% of his time trying to prove something which is false,” effort that could be avoided with technology that finds a counterexample automatically, or that finds the smallest counterexample. “That gives so much insight that it can really help mathematicians do their work faster.” **C**

Further Reading

Hartnett, K.

Computer Search Settles 90-Year-Old Math Problem, *Quanta*, August 2020, <https://bit.ly/2Xh1ni4>

Brakensiek, J., Heule, M.J.H., Mackey, J., and Narváez, D. The Resolution of Keller’s Conjecture <http://bit.ly/3olSXPB>

Hales, T., et al.

A Formal Proof of the Kepler Conjecture, *Forum of Mathematics, Pi* 5, e2 (2017). DOI: <https://doi.org/10.1017/fmp.2017.1>

Lagarias, J.C., and Shor, P.W.

Keller’s cube-tiling conjecture is false in high dimensions *Bull. Amer. Math. Soc.* 27, 279-283. (1992). DOI: <https://doi.org/10.1090/S0273-0979-1992-00318-X>

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

© 2021 ACM 0001-0782/21/5 \$15.00