Science | DOI:10.1145/3374874 Chris Edwards

Learning to Trust Quantum Computers

They need to show us they can solve the biggest problems.

One of the core beliefs behind the push to build quantum computers is that they will power a massive expansion in computing capability. However, how much capability could the technology really bring and, even if we can harness all that power, how can we be sure quantum computing will provide accurate answers when there is no way to run the same algorithms on conventional computers for verification?

A paper on the use of quantum entanglement in verifying the solutions to problems published in the spring of 2019 by California Institute of Technology (Caltech) postdoctoral researchers Anand Natarajan and John Wright has shown how quantum computers can prove their results are legitimate. The expansion in what is provable is likely to lead to a situation where the ability of quantum computers to demonstrate the correctness of their calculations far outstrips their ability to compute the results in the first place.

The key to checking the work of highly powerful computers lies in a result published in 1988 by Michael Ben-Or and Avi Wigderson of Hebrew University, working together with Shafi Goldwasser and Joe Kilian at the Massachusetts

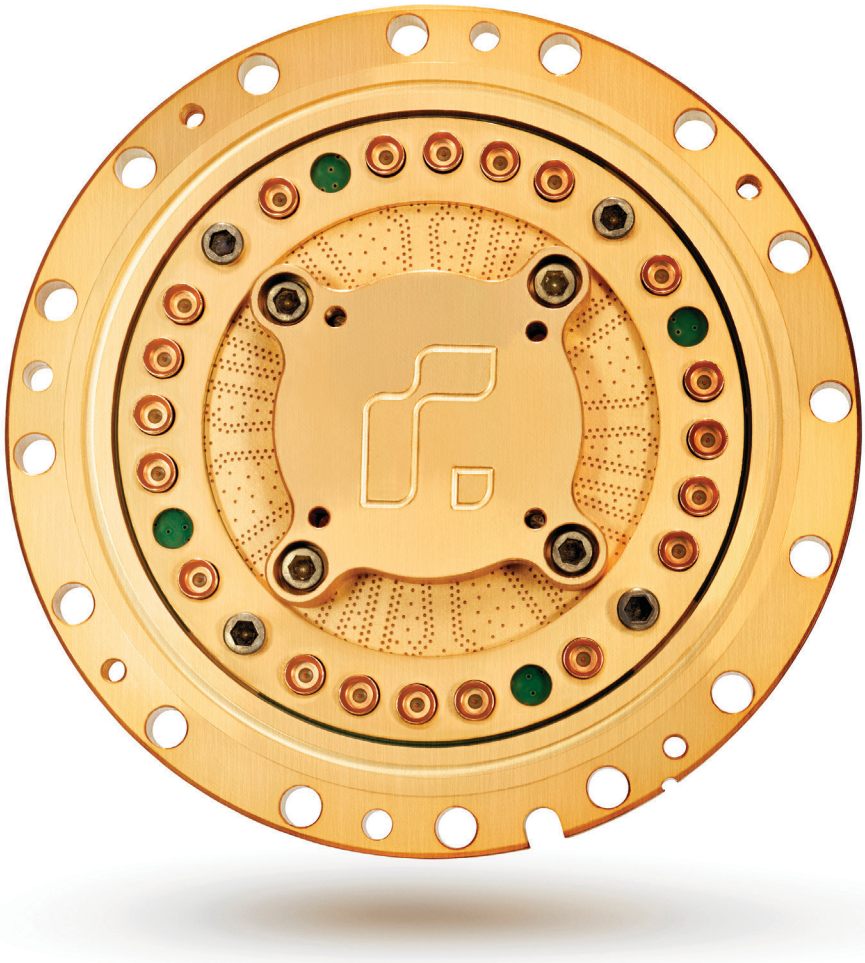

A Rigetti Computing quantum processor based on 32-qubit superconducting chip technology.

IMAGE COURTESY RIGETTI COMPUTING. PHOTO BY JUSTIN FANTL.

Institute of Technology (MIT). Their original aim was to find a new way to construct authentication systems that did not rely on cryptographic functions that are assumed to be hard for computers to break. To do so, they employed the idea of the zero-knowledge proof, which lets two remote systems (known as provers) demonstrate to a third party, the verifier, that they hold a secret, without revealing the secret itself.

The provers are assumed to be able to solve any problem. To prevent the provers from cheating, the verifier uses randomly generated queries in an interactive protocol designed to catch attempts by either prover to conspire with the other to deliver a false answer. The only way they could lie reliably is to work together. To avoid this possibility, classical implementations of zero-knowledge proofs rely on setting up the tests so the two provers cannot communicate with each other directly.

The work on multiprover interactive proof (MIP) systems later expanded into delegated computing: to let systems offload tasks to remote servers. In a paper presented at the 2008 ACM Symposium of the Theory of Computing in British Columbia, Canada, Goldwasser and colleagues referenced the Harry Potter series, whose seventh and final volume was published that year, with the claim that a "muggle" machine (in this context, a conventional computer) could check the work much more capable computers claim to be able to perform. The question that challenged theoreticians was how much more those magical systems could do, and still let non-wizardly computers check whether their work is valid.

A few years after the original MIP work appeared, László Babai, Lance Fortnow, and Carsten Lund, working at the University of Chicago, showed a user with access to a machine able only to handle problems in polynomial time could verify the work on problems that, for a Turing machine that can perform multiple operations in each computational step, require exponential time to solve. This is a complexity class known as NEXP. The open question before last year was what difference quantum entanglement between the provers would make, an arrangement that mathematicians call MIP*.

"When MIP* was first introduced, it wasn't clear whether it was more powerful than MIP," Wright says.

Researchers believed MIP* could deliver greater power by letting the provers share states through entanglement. But it carried with it the threat of collusion. Separate provers could use their shared knowledge to collude with each other in a way that classical multiprover systems do not. This would, in turn, reduce the theoretical power of multiprover systems that use quantum-capable provers.

Milestones

ACM Recognizes 2019 Fellows

ACM has named 58 members ACM Fellows for their wide-ranging, fundamental contributions in areas including artificial intelligence, cloud computing, combating cybercrime, quantum computing, and wireless networking.

"Computing technology has had a tremendous impact in shaping how we live and work today," said ACM President Cherri M. Pancake. "In highlighting the accomplishments of the ACM Fellows, we hope to give credit where it is due, while also educating the public about the extraordinary array of areas in which computing professionals work."

The 2019 Fellows hail from universities, companies and research centers in Australia, Canada, China, Egypt, France, Germany, Israel, Italy, Switzerland, and the United States.

The 2019 ACM Fellows are:

Scott J. Aaronson, ***University of Texas***
Tarek F. Abdelzaher, ***University of Illinois at Urbana-Champaign***
Saman Amarasinghe, ***Massachusetts Institute of Technology***
Kavita Bala, ***Cornell University***
Magdalena Balazinska, ***University of Washington***
Paul Beame, ***University of Washington***
Emery D. Berger, ***University of Massachusetts Amherst***
Ronald F. Boisvert, ***National Institute of Standards and Technology***
Christian Cachin, ***University of Bern***
Brad Calder, ***Google***
Diego Calvanese, ***Free University of Bozen-Bolzano***
Srdjan Capkun, ***Swiss Federal Polytechnic, Zurich***
Claire Cardie, ***Cornell University***
Timothy M. Chan, ***University of Illinois at Urbana-Champaign***
Kanianthra Mani Chandy, ***California Institute of Technology***
Xilin Chen, ***Institute of Computing Technology, Chinese Academy of Sciences***
Elizabeth F. Churchill, ***Google***
Philip R. Cohen, ***Monash University***
Vincent Conitzer, ***Duke University***
Noshir Contractor, ***Northwestern University***
Matthew B. Dwyer, ***University of Virginia***
Elena Ferrari, ***University of Insubria***
Michael J. Freedman, ***Princeton University***
Deborah Frincke, ***U.S. National Security Agency***
Lise Getoor, ***University of California, Santa Cruz***
Maria L. Gini, ***University of Minnesota***
Subbarao Kambhampati, ***Arizona State University***
Tamara G. Kolda, ***Sandia National Laboratories***
Songwu Lu, ***University of California, Los Angeles***
Wendy Elizabeth Mackay, ***Inria***
Diana Marculescu, ***University of Texas at Austin***
Sheila McIlraith, ***University of Toronto***
Rada Mihalcea, ***University of Michigan***
Robin R. Murphy, ***Texas A&M University***
Marc Najork, ***Google***
Jason Nieh, ***Columbia University***
Hanspeter Pfister, ***Harvard University***
Timothy M. Pinkston, ***University of Southern California***
Mihai Pop, ***University of Maryland, College Park***
Andreas Reuter, ***Heidelberg University/Heidelberg Laureate Forum Foundation***
Jeffrey S. Rosenschein, ***Hebrew University***
Srinivasan Seshan, ***Carnegie Mellon University***
Prashant J. Shenoy, ***University of Massachusetts Amherst***
Peter W. Shor, ***Massachusetts Institute of Technology***
Mona Singh, ***Princeton University***
Ramesh K. Sitaraman, ***University of Massachusetts Amherst***
Dawn Song, ***University of California, Berkeley***
Salvatore J. Stolfo, ***Columbia University***
Dacheng Tao, ***The University of Sydney***
Moshe Tennenholtz, ***Technion***
Giovanni Vigna, ***University of California, Santa Barbara***
Nisheeth K. Vishnoi. ***Yale University***
Darrell Whitley, ***Colorado State University***
Yuan Xie, ***University of California, Santa Barbara***
Moustafa Amin Youssef, ***Alexandria University***
Carlo A. Zaniolo. ***University of California, Los Angeles***
Lidong Zhou, ***Microsoft Research Asia***

ACM will formally recognize its 2019 Fellows at the annual Awards Banquet in San Francisco on June 20, 2020. Additional information about the 2019 ACM Fellows, as well as previously named ACM Fellows, is available on the ACM Fellows site at https://awards.acm.org/fellows.

Working with Caltech professor of computing and mathematical sciences Thomas Vidick, Natarajan demonstrated in 2016 it was possible to stop cheating by entangled provers. The protocol, known as the Pauli braiding test, exploits the Heisenberg Uncertainty Principle. If one system attempts to measure a variable held by a quantum bit or qubit, information on the other properties the entangled qubit holds is destroyed. Over a series of questions, the verifier switches between the provers, asking them different things in a way that the answers can be checked for consistency, but destroying data that would let them collude.

By this point, mathematicians in the field considered the quantum versions of multiprover systems at least as powerful as their entirely classical counterparts. What remains unclear to this day is the upper bound. Natarajan and Wright were able to expand the known upper bound to systems exponentially larger than NEXP: a class called NEEXP.

A major hurdle in developing the proof was the mismatch between the size of messages a verifier can send to the provers and the space occupied by a problem with an NEEXP-level of complexity. Simply expressing the problem can exhaust the capacity of the verifier. "The graph becomes so large that to even give a name to a specific vertex requires an exponential number of bits," Wright says.

If a verifier, for example, cannot identify a specific vertex in the complete graph, how does it get provers to deliver believable answers? The answer, for Natarajan and Wright, was to use what they call "introspection."

Says Henry Yuen, assistant professor of computer science and mathematics at the University of Toronto, "The idea can be traced back to a paper of Zhengfeng Ji, although he didn't call it introspection."

Ji's work led to a result superficially similar to Natarajan and Wright's, but with a key difference: there would be a credibility gap in what the provers can demonstrate to the verifier that expands as problems become bigger. Ji also showed a carefully chosen protocol can make provers ask questions of themselves in a controlled manner and use the answers they derive to look into the much, much larger answer space. In effect, introspection acts as a form of compression that makes it possible for classical machines to handle spaces far beyond their capabilities.

The problem space that a MIP* system can handle could be far bigger than that and introspection may provide the way forward. "You'd hope to use successive layers of introspection," Wright says, with each round of introspection producing a smaller protocol until the messages are small enough to be exchanged with a purely classical computer. Here again, Ji's work provides inspiration for this nested use of introspection.

In a follow-up paper with Yuen, Vidick, and Singapore University of Technology and Design associate professor Joseph Fitzsimons, Ji showed how it was possible to use the technique recursively to arbitrarily deep levels of exponentiation. As long as the verifier could tolerate some level of doubt, they could check the work of provers on problem spaces far larger than the number of atoms in the universe. For mathematicians specializing in the field, Natarajan and Wright's result put MIP* on much firmer footing that may expand the power of quantum-enabled provers to the edges of complexity theory.

"This result, I think, has dramatically shifted the attitude of the community towards MIP*: while before I think many of us would've hedged our bets on the complexity of MIP*, now it seems much more conceivable—likely, even—that the complexity of MIP* won't just stop at NEEXP, but could keep on going to arbitrarily large complexity classes such as NEEEXP or NEEEEEEEXP, or potentially even undecidable problems," Yuen says. "The Natarajan-Wright result is giving us really compelling evidence of the tremendous complexity of MIP*."

Wright says work continues on recursively using introspection, though it is not yet clear whether it will work. The implications would be pretty big, he says, if undecidable problems are also shown to fit into MIP*. "If two people told you that a Turing machine halts on a given input, how would you check that without just running it until it halts, which could take an unbounded amount of time?" he asks. Quantum multiprovers, in this scenario, would provide the way to determine whether those people were telling the truth. How they knew would likely remain a mystery as there is no prospect of anyone being able to construct a computer that can provide the answer.

The result also circles back to the zero-knowledge proofs of cryptography research that kicked off work on multiprover systems. Says Yuen, "One of my recent papers shows that the class MIP* is equal to the class zero-knowledge MIP*. In other words, every interactive proof conducted with quantum-entangled provers can be transformed into an equivalent zero-knowledge protocol."

If undecidable problems do fall into MIP*, the verifier would have proof there is an answer to a problem that computer science today considers unknowable, but not what the answer is. "While I can't imagine that there could be practical applications of this fact, it would be an amusing fact if nothing else."

Yuen says there are potential practical benefits for much smaller problems: "One could envision people inventing protocols for verifying extremely large computations where the verifier could be extremely succinct in its interrogation of the provers."

The research continues, although much remains uncertain as to how far mathematicians can push the upper bounds of theoretical capability. What has been shown so far is that quantum entanglement could underpin a massive expansion in what is computable.

Further Reading

Natarajan, A, & Wright J.
NEEXP ⊆ MIP*
arXiv preprint (2019). arXiv: 1904.05870

Goldwasser, S., Kalai, Y.T., and Rothblum, G.N.
Delegating Computation: Interactive Proofs for Muggles
Proceedings of the 40th Annual Symposium on the Theory of Computing (STOC), May 2008), DOI: 10.1145/1374376.1374396

Ji, Z.
Compression of Quantum Multiprover Interactive Proofs
arXiv preprint (2016). arXiv: 1610.03133

Grilo, A.B, Slofstra, W., and Yuen, H.
Perfect Zero Knowledge for Quantum Multiprover Interactive Proofs
arXiv preprint (2019). arXiv: 1905.11280

Chris Edwards is a Surrey, U.K.-based writer who reports on electronics, IT, and synthetic biology.