Graduate Student Solves Quantum Verification Problem

Urmila Mahadev spent eight years in graduate school solving one of the most basic questions in quantum computation: How do you know whether a quantum computer has done anything quantum at all?

By Erica Klarreich

Jana Ašenbrennerová for Quanta Magazine

Urmila Mahadev giving a computer science seminar last week at the University of California, Berkeley, ahead of her presentation yesterday at the Symposium on Foundations of Computer Science in Paris.

In the spring of 2017, Urmila Mahadev found herself in what most graduate students would consider a pretty sweet position. She had just solved a major problem in quantum computation, the study of computers that derive their power from the strange laws of quantum physics. Combined with her

earlier papers, Mahadev's new result, on what is called blind computation, made it "clear she was a rising star," said Scott Aaronson, a computer scientist at the University of Texas, Austin.

Mahadev, who was 28 at the time, was already in her seventh year of graduate school at the University of California, Berkeley — long past the stage when most students become impatient to graduate. Now, finally, she had the makings of a "very beautiful Ph.D. dissertation," said Umesh Vazirani, her doctoral adviser at Berkeley.

But Mahadev did not graduate that year. She didn't even consider graduating. She wasn't finished.

For more than five years, she'd had a different research problem in her sights, one that Aaronson called "one of the most basic questions you can ask in quantum computation." Namely: If you ask a quantum computer to perform a computation for you, how can you know whether it has really followed your instructions, or even done anything quantum at all?

This question may soon be far from academic. Before too many years have elapsed, researchers hope, quantum computers may be able to offer exponential speedups on a host of problems, from modeling the behavior around a black hole to simulating how a large protein folds up.

But once a quantum computer can perform computations a classical computer can't, how will we know if it has done them correctly? If you distrust an ordinary computer, you can, in theory, scrutinize every step of its computations for yourself. But quantum systems are fundamentally resistant to this kind of checking. For one thing, their inner workings are incredibly complex: Writing down a description of the internal state of a computer with just a few hundred quantum bits (or "qubits") would require a hard drive larger than the entire visible universe.

And even if you somehow had enough space to write down this description, there would be no way to get at it. The inner state of a quantum computer is generally a superposition of many different non-quantum, "classical" states (like Schrödinger's cat, which is simultaneously dead and alive). But as soon as you measure a quantum state, it collapses into just one of these classical states. Peer inside a 300-qubit quantum computer, and essentially all you will see is 300 classical bits — zeros and ones — smiling blandly up at you.

"A quantum computer is very powerful, but it's also very secretive," Vazirani said.

Given these constraints, computer scientists have long wondered whether it is possible for a quantum computer to provide any ironclad guarantee that it really has done what it claimed. "Is the interaction between the quantum and the classical worlds strong enough so that a dialogue is possible?" asked Dorit Aharonov, a computer scientist at the Hebrew University of Jerusalem.

During her second year of graduate school, Mahadev became captivated by this problem, for reasons even she doesn't fully understand. In the years that followed, she tried one approach after another. "I've had a lot of moments where I think I'm doing things right, and then they break, either very quickly or after a year," she said.

But she refused to give up. Mahadev displayed a level of sustained determination that Vazirani has never seen matched. "Urmila is just absolutely extraordinary in this sense," he said.

Jana Ašenbrennerová for Quanta Magazine

Now, after eight years of graduate school, Mahadev has succeeded. She has come up with an interactive protocol by which users with no quantum powers of their own can nevertheless employ cryptography to put a harness on a quantum computer and drive it wherever they want, with the certainty that the quantum computer is following their orders. Mahadev's approach, Vazirani said, gives the user "leverage that the computer just can't shake off."

For a graduate student to achieve such a result as a solo effort is "pretty astounding," Aaronson said.

Mahadev, who is now a postdoctoral researcher at Berkeley, presented her protocol yesterday at the annual Symposium on Foundations of Computer Science, one of theoretical computer science's biggest conferences, held this year in Paris. Her work has been awarded the meeting's "best paper" and "best student paper" prizes, a rare honor for a theoretical computer scientist.

In a blog post, Thomas Vidick, a computer scientist at the California Institute of Technology who has collaborated with Mahadev in the past, called her result "one of the most outstanding ideas to have emerged at the interface of quantum computing and theoretical computer science in recent years."

Quantum computation researchers are excited not just about what Mahadev's protocol achieves, but also about the radically new approach she has brought to bear on the problem. Using classical cryptography in the quantum realm is a "truly novel idea," Vidick wrote. "I expect many more results to continue building on these ideas."

A Long Road

Raised in Los Angeles in a family of doctors, Mahadev attended the University of Southern California, where she wandered from one area of study to another, at first convinced only that she did not want to become a doctor herself. Then a class taught by the computer scientist Leonard Adleman, one of the creators of the famous RSA encryption algorithm, got her excited about theoretical computer science. She applied to graduate school at Berkeley, explaining in her application that she was interested in all aspects of theoretical computer science — except for

quantum computation.

“It sounded like the most foreign thing, the thing I knew least about,” she said.

But once she was at Berkeley, Vazirani’s accessible explanations soon changed her mind. He introduced her to the question of finding a protocol for verifying a quantum computation, and the problem “really fired up her imagination,” Vazirani said.

“Protocols are like puzzles,” Mahadev explained. “To me, they seem easier to get into than other questions, because you can immediately start thinking of protocols yourself and then breaking them, and that lets you see how they work.” She chose the problem for her doctoral research, launching herself on what Vazirani called “a very long road.”

If a quantum computer can solve a problem that a classical computer cannot, that doesn’t automatically mean the solution will be hard to check. Take, for example, the problem of factoring large numbers, a task that a big quantum computer could solve efficiently, but which is thought to be beyond the reach of any classical computer. Even if a classical computer can’t factor a number, it can easily check whether a quantum computer’s factorization is correct — it just needs to multiply the factors together and see if they produce the right answer.

Yet computer scientists believe (and have recently taken a step toward proving) that many of the problems a quantum computer could solve do not have this feature. In other words, a classical computer not only cannot solve them, but cannot even recognize whether a proposed solution is correct. In light of this, around 2004, Daniel Gottesman — a physicist at the Perimeter Institute for Theoretical Physics in Waterloo, Ontario — posed the question of whether it is possible to come up with any protocol by which a quantum computer can prove to a non-quantum observer that it really has done what it claimed.

Jana Ašenbrennerová for Quanta Magazine

Within four years, quantum computation researchers had achieved a partial answer. It is possible, two different teams showed, for a quantum computer to prove its computations, not to a purely classical verifier, but to a verifier who has access to a very small quantum computer of her own. Researchers later refined this approach to show that all the verifier needs is the capacity to measure a single qubit at a time.

And in 2012, a team of researchers including Vazirani showed that a completely classical verifier could check quantum computations if they were carried out by a pair of quantum computers that can't communicate with each other. But that paper's approach was tailored to this specific scenario, and the problem seemed to hit a dead end there, Gottesman said. "I think there were probably people who thought you couldn't go further."

It was around this time that Mahadev encountered the verification problem. At first, she tried to come up with an "unconditional" result, one that makes no assumptions about what a quantum computer can or cannot do. But after she had worked on the problem for a while with no progress, Vazirani proposed instead the possibility of using "post-quantum" cryptography — that is, cryptography that researchers believe is beyond the capability of even a quantum computer to break, although they don't know for sure. (Methods such as the RSA algorithm that are used to encrypt things like online transactions are not post-quantum — a large quantum computer could break them, because their security depends on the hardness of factoring large numbers.)

In 2016, while working on a different problem, Mahadev and Vazirani made an advance that would later prove crucial. In collaboration with Paul Christiano, a computer scientist now at OpenAI, a company in San Francisco, they developed a way to use cryptography to get a quantum computer to build what we'll call a "secret state" — one whose description is known to the classical verifier, but not to the quantum computer itself.

Their procedure relies on what's called a "trapdoor" function — one that is easy to carry out, but hard to reverse unless you possess a secret cryptographic key. (The researchers didn't know how to actually build a suitable trapdoor function yet — that would come later.) The function is also required to be "two-to-one," meaning that every output corresponds to two different inputs. Think, for example of the function that squares numbers — apart from the number 0, each output (such as 9) has two corresponding inputs (3 and −3).

Armed with such a function, you can get a quantum computer to create a secret state as follows: First, you ask the computer to build a superposition of all the possible inputs to the function (this might sound complicated for the computer to carry out, but it's actually easy). Then, you tell the computer to apply the function to this giant superposition, creating a new state that is a superposition of all the possible outputs of the function. The input and output superpositions will be entangled, which means that a measurement on one of them will instantly affect the other.

Next, you ask the computer to measure the output state and tell you the result. This measurement collapses the output state down to just one of the possible outputs, and the input state instantly collapses to match it, since they are entangled — for instance, if you use the squaring function, then if the output is the 9 state, the input will collapse down to a superposition of the 3 and −3 states.

But remember that you're using a trapdoor function. You have the trapdoor's secret key, so you can easily figure out the two states that make up the input superposition. But the quantum computer cannot. And it can't simply measure the input superposition to figure out what it is made of, because

that measurement would collapse it further, leaving the computer with one of the two inputs but no way to figure out the other.

In 2017, Mahadev figured out how to build the trapdoor functions at the core of the secret-state method by using a type of cryptography called Learning With Errors (LWE). Using these trapdoor functions, she was able to create a quantum version of "blind" computation, by which cloud-computing users can mask their data so the cloud computer can't read it, even while it is computing on it. And shortly after that, Mahadev, Vazirani and Christiano teamed up with Vidick and Zvika Brakerski (of the Weizmann Institute of Science in Israel) to refine these trapdoor functions still further, using the secret-state method to develop a foolproof way for a quantum computer to generate provably random numbers.

Mahadev could have graduated on the strength of these results, but she was determined to keep working until she had solved the verification problem. "I was never thinking of graduation, because my goal was never graduation," she said.

Not knowing whether she would be able to solve it was stressful at times. But, she said, "I was spending time learning about things that I was interested in, so it couldn't really be a waste of time."

Set in Stone

Mahadev tried various ways of getting from the secret-state method to a verification protocol, but for a while she got nowhere. Then she had a thought: Researchers had already shown that a verifier can check a quantum computer if the verifier is capable of measuring quantum bits. A classical verifier lacks this capability, by definition. But what if the classical verifier could somehow force the quantum computer to perform the measurements itself and report them honestly?

The tricky part, Mahadev realized, would be to get the quantum computer to commit to which state it was going to measure before it knew which kind of measurement the verifier would ask for — otherwise, it would be easy for the computer to fool the verifier. That's where the secret-state method comes into play: Mahadev's protocol requires the quantum computer to first create a secret state and then entangle it with the state it is supposed to measure. Only then does the computer find out what kind of measurement to perform.

Since the computer doesn't know the makeup of the secret state but the verifier does, Mahadev showed that it's impossible for the quantum computer to cheat significantly without leaving unmistakable traces of its duplicity. Essentially, Vidick wrote, the qubits the computer is to measure have been "set in cryptographic stone." Because of this, if the measurement results look like a correct proof, the verifier can feel confident that they really are.

"It is such a wonderful idea!" Vidick wrote. "It stuns me every time Urmila explains it."

Mahadev's verification protocol — along with the random-number generator and the blind encryption method — depends on the assumption that quantum computers cannot crack LWE. At present, LWE is widely regarded as a leading candidate for post-quantum cryptography, and it may soon be adopted by the National Institute of Standards and Technology as its new cryptographic standard, to replace the ones a quantum computer could break. That doesn't guarantee that it really is secure against quantum computers, Gottesman cautioned. "But so far it's solid," he said. "No one has found evidence that it's likely to be breakable."

In any case, the protocol's reliance on LWE gives Mahadev's work a win-win flavor, Vidick wrote. The only way that a quantum computer could fool the protocol is if someone in the quantum

computing world figured out how to break LWE, which would itself be a remarkable achievement.

Mahadev's protocol is unlikely to be implemented in a real quantum computer in the immediate future. For the time being, the protocol requires too much computing power to be practical. But that could change in the coming years, as quantum computers get larger and researchers streamline the protocol.

Mahadev's protocol probably won't be feasible within, say, the next five years, but "it is not completely off in fantasyland either," Aaronson said. "It is something you could start thinking about, if all goes well, at the next stage of the evolution of quantum computers."

And given how quickly the field is now moving, that stage could arrive sooner rather than later. After all, just five years ago, Vidick said, researchers thought that it would be many years before a quantum computer could solve any problem that a classical computer cannot. "Now," he said, "people think it's going to happen in a year or two."

As for Mahadev, solving her favorite problem has left her feeling a bit at sea. She wishes she could understand just what it was about that problem that made it right for her, she said. "I have to find a new question now, so it would be nice to know."

But theoretical computer scientists see Mahadev's unification of quantum computation and cryptography not so much as the end of a story, but as the initial exploration of what will hopefully prove a rich vein of ideas.

"My feeling is that there are going to be lots of follow-ups," Aharonov said. "I'm looking forward to more results from Urmila."