COMPUTING

Inside the quest for unbreakable encryption

Cryptographers want encryption schemes that are impossible for tomorrow's quantum computers to crack. There's only one catch: they might not exist.

By Stephen Ornes

October 19, 2023

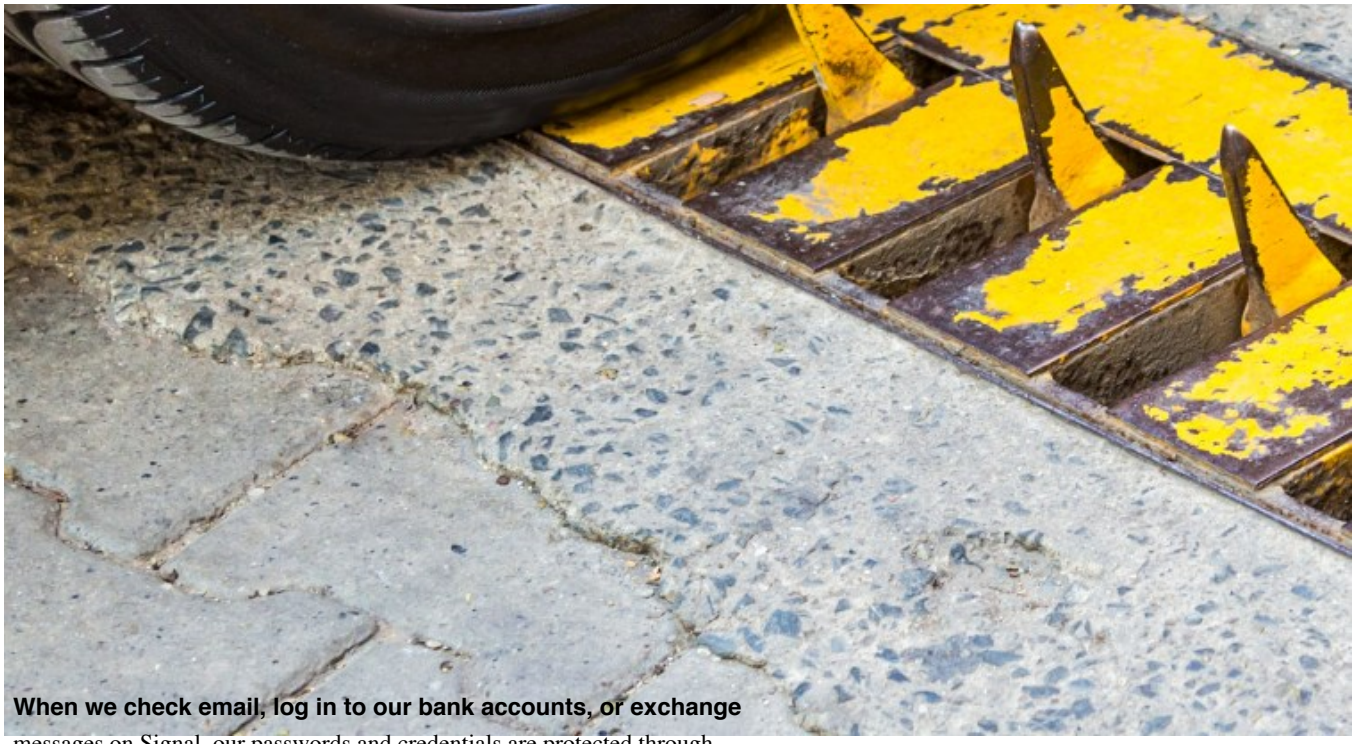

When we check email, log in to our bank accounts, or exchange messages on Signal, our passwords and credentials are protected through encryption, a locking scheme that uses secrets to disguise our data. It works like a cyber padlock: with the right key someone can unlock the data. Without it, they'll have to resort to laborious brute-force methods, the digital equivalent of hacksaws and blowtorches.

GETTY IMAGES

Our trust in online security is rooted in mathematics. Encryption schemes are built on families of math problems called one-way functions—calculations that are easy to carry out in one direction but almost impossible to solve efficiently from the other, even with a powerful computer. They're sort of a computational equivalent of those road spikes found at the exits of airport car rental agencies. Drive in one direction and you barely notice. Hit reverse and you won't get far (and will need new tires).

There's a problem, however. Although mathematicians suspect true one-way functions exist, they have yet to prove it. They haven't proved that the thorny problems we *do* use are impossible, or even extremely impractical, to solve. Instead, it could just be that we haven't yet found the appropriate mathematical means to take the problems apart. This conundrum haunts all encryption. Our data is secured by the fact that no one knows how to crack the schemes that protect it—at least not yet.

It's not just today's hackers we may need to worry about. Security experts have long warned of a threat that hasn't yet materialized: quantum computers. In the future these machines could execute a program that quickly solves the math problems behind today's state-of-the-art encryption. That threat puts personal financial, medical, and other information at risk. Hackers could steal today's encrypted data and store it away, just waiting for the arrival of new technological lockpicks.

Computer scientists, mathematicians, and cryptographers are on a quest to find new

POPULAR

This new data poisoning too
fight back against generative
Melissa Heikkilä

DeepMind's cofounder: Gen
just a phase. What's next is
Will Douglas Heaven

Everything you need to knov
artificial wombs
Cassandra Willyard

Deepfakes of Chinese influe
livestreaming 24/7
Zeyi Yang

encryption algorithms that can withstand attacks not only from today's conventional computers but also from tomorrow's quantum machines. What they want is a big, sticky math problem—something that's robust enough to withstand attacks from classical and quantum computers but can still be easily implemented in cyberspace.

Unfortunately, no one has yet found a single type of problem that is provably hard for computers—classical or quantum—to solve. (In the world of cryptography, "hard" describes a problem whose solution requires an unreasonable number of steps or amount of computing power.) If one-way functions don't exist, then cryptographers' whack-a-mole process of finding flaws and developing ever stronger schemes to block clever hackers will persist indefinitely.

"The question of whether one-way functions exist is really the most important problem," says Rafael Pass, a theoretical computer scientist at Tel Aviv University in Israel. It's a conundrum that dates to the 1970s and the dawn of a research area now known as computational complexity theory. Over five decades, theorists and cryptographers have been looking for ways to establish whether such functions do exist. Perhaps the problems we hope or suspect are one-way are just easier, breakable ones in disguise.

Pass is exploring how one-way functions are connected to a raft of other open problems, a promising line of research that has drawn other theorists into the quest. At the same time, people focused on the practical side of cryptography are plowing ahead, hunting for new schemes that are—if not provably hard—seemingly strong enough to hold up against quantum computers.

Computer scientists find themselves at a curious crossroads, unsure of whether post-quantum algorithms are truly unassailable—or just believed to be so.

For the last seven years, the job of finding the best candidates has been spearheaded by the National Institute of Standards and Technology (NIST), the US government body charged with collecting, testing, and standardizing cryptographic algorithms for public use. NIST has been running dozens of potential "post-quantum" algorithms through a gauntlet of tests and making them available for outside testing. The process has winnowed the field to a few finalists, and in August NIST announced that one called CRYSTALS-Kyber, which takes an approach believed to be robust enough to counter quantum attacks, will be the first to be officially recommended for public use by 2024. After that, companies and governments will adopt the algorithm for encrypting data.

Will it hold up? The answer will help determine the trajectory of cybersecurity in the near term. But it's far from settled: history suggests that our faith in unbreakability has often been misplaced, and over the years, seemingly impenetrable encryption candidates have fallen to surprisingly simple attacks. Computer scientists find themselves at a curious crossroads, unsure of whether post-quantum algorithms are truly unassailable—or just believed to be so. It's a distinction at the heart of modern encryption security.

The myth and reality of unbreakability

Securing secret messages hasn't always been tied to difficult math problems; until

recently, cryptography was barely mathematical at all. In ancient Greece, military leaders encoded messages using a scytale, a cylindrical device that revealed a hidden message when a strip of seemingly jumbled text was wound around it. Centuries later, Roman historians described a code, often attributed to Julius Caesar, that involved shifting letters in a message three spots up in the alphabet; for example, a *d* would be written as an *a*.

Related Story

What's next for quantum computing

Companies are moving away from setting qubit records in favor of practical hardware and long-term goals.

In history as in our modern world, secret codes were frequently broken. In the 16th century, during the decades she spent imprisoned by her cousin Queen Elizabeth I, Mary, Queen of Scots, used elaborate, symbol-based ciphers to encode hundreds of letters, most of which were aimed at securing her freedom and regaining the throne. She didn't prevail: Elizabeth I's team of spies and codebreakers intercepted, decoded, and copied the letters. In the one that sealed her fate, Mary approved of a plan to assassinate Elizabeth with six words: "sett the six gentlemen to woork." In response, Elizabeth eventually ordered her cousin beheaded in 1587.

In 1932, codebreakers in Poland cracked the code for Germany's early Enigma machine, invented at the end of World War I. They later shared their intel with British codebreakers, who cracked a more advanced version of Enigma during World War II.

Pass, the theoretical computer scientist in Tel Aviv, half-jokingly refers to all time before the 1970s as the "dark age of cryptography."

"Cryptography wasn't really a scientific field," he says. "It was more like artist versus attackers. You needed to have [artistic] skills to invent an encryption scheme. And then it would get deployed until some clever person would figure out how to break it. And it was just going on and on like that."

That changed, Pass says, in November 1976, when cryptographers Whitfield Diffie and Martin Hellman, at Stanford, described a novel way for two people to devise a key that only they knew—one they could then use to pass secret messages. Crucially, they wouldn't have to meet to do it. This was a groundbreaking notion. Previously, both sender and receiver had to physically possess a key for encoding and decoding. To decrypt a message encoded with the Enigma machine, for example, a recipient needed a key sheet that revealed the initial encryption settings.

The secret to the Diffie-Hellman strategy was for two people to build the key using a straightforward mathematical problem that's easy to compute in one direction and laborious in the other. Here's how it works: The two people who want to communicate secretly, usually designated Alice and Bob in these setups, each pick a secret number. Then, together, they agree on a pair of numbers that they share publicly (one is a big prime, and the other is called the base). Each of them next carries out a series of mathematical operations to combine those private numbers with the prime and the base.

Then they exchange the results, and they each carry out another series of mathematical operations on the new numbers. In the end, both Alice and Bob will

have done the same operations on the same numbers—just not in the same order—and arrived at the same answer. The digits of that answer become the encryption. And an eavesdropper who intercepts the transmission—often nicknamed Eve—won't be able to easily unravel the mathematical jumble without knowing at least one of the private numbers. She could start testing numbers in a brute-force approach, but that would require an unreasonable amount of calculation.

The complicated problem that Eve would have to solve is called finding a discrete logarithm. The Diffie-Hellman approach is still used today—to secure some VPNs, for example—and is integral to some post-quantum schemes.

In their paper, Diffie and Hellman noted that there was no existing algorithm capable of solving the discrete log problem in a reasonable amount of time. There still isn't. They went on to introduce, for the first time, the notion of one-way functions as a basis for secure cryptography.

Today, secure online interactions that involve authentication or digital signatures, for example, are based on that general idea. But without mathematical proof that the problems they rely on are one-way functions, the possibility remains that someone might discover an efficient scheme for cracking them.

The quantum menace

Today, online transactions begin with a kind of digital handshake, and the security of that handshake is often guaranteed by another math problem that's presumed to be difficult. The most popular encryption scheme used today was introduced in 1977 by a trio of young computer scientists who were energized by Diffie and Hellman's 1976 paper. They called their approach RSA, after the last names of the scientists (Ron Rivest, Adi Shamir, and Leonard Adleman).

RSA, which is based on the difficulty of finding prime factors relative to the ease of multiplying them together, is a bit different from the Diffie-Hellman approach. Diffie-Hellman is a shared secret: it allows two users to devise a key over an insecure channel (like the internet), and that key is used to disguise messages. In RSA, Alice uses Bob's key—based on big prime numbers—to encrypt a message that only he can unlock. RSA can secure data sent from one person to another.

It quickly became one of the most popular public-key encryption methods. It's easy to use and adapt. Over time, as new algorithms have emerged that can factor faster, and computers have become more powerful, NIST has recommended using larger and larger numbers for security. The numbers are represented in binary form with *1*s and *0*s, and these binary digits are better known as "bits." The number 13, for example, is written in binary as *1101*, which has four bits. NIST currently recommends using a key represented by at least 2,048 bits—which corresponds to a number with over 600 digits. (To date, the largest number that has been factored into two primes was made up of 250 digits, and the process took nearly 3,000 hours of computing time.) That's a strength of RSA—even if it's not uncrackable, it's been easy to keep upping the ante, making it computationally impractical to break.

In 1994, however, a threat of a different type emerged when the American mathematician Peter Shor, then at Bell Labs, devised an algorithm for quantum computers that could solve the factoring problem in a reasonable amount of time.

(It was a double threat: his approach could also conquer the discrete log problem in the Diffie-Hellman approach.)

Shor's paper ignited excitement and anxiety among those who wanted to build quantum computers and those who recognized the threat it posed to cybersecurity. Fortunately for cryptographers, not just any quantum computer would do.

The last three decades of cybersecurity have played out like an increasingly intricate game, with researchers perpetually building and breaking—or attempting to break—new candidates.

A few years back, researchers at Google and the KTH Royal Institute of Technology, in Sweden, estimated that it would take a quantum computer composed of 20 million quantum bits, or qubits, some eight hours to break today's 2,048-bit RSA security. Current state-of-the-art machines are nowhere close to that size: the largest quantum computer to date, built by IBM, debuted last year with 433 qubits.

Whether or not RSA can be considered at immediate risk of a quantum attack depends largely on whom you ask, says computer scientist Ted Shorter, who cofounded the cybersecurity company Keyfactor. He sees a cultural divide between the theorists who study the mathematics of encryption and the cryptographers who work in implementation.

To some, the end seems nigh. "You talk to a theoretical computer scientist and they're like, *Yes, RSA is done*, because they can imagine it," Shorter says. For them, he adds, the existence of Shor's algorithm points to the end of encryption as we know it.

Many cryptographers who are implementing real-world security systems are less concerned about the quantum future than they are about today's cleverest hackers. After all, people have been trying to factor efficiently for thousands of years, and now the only known method requires a computer that doesn't exist.

Thomas Decru, a cryptographer at KU Leuven in Belgium, says the quantum threat must be taken seriously, but it's hard to know if RSA will fall to quantum computers in five years or longer—or never. "As long as quantum computers do not exist, everything you say about them is speculative, in a way," he says. Pass is more certain about the threat: "It's safe to say that the existence of this quantum algorithm means there are cracks in the problem, right?"

The thorns of implementation

But we have to be ready for anything, says Lily Chen, a mathematician who manages NIST's Cryptographic Technology Group and works on the ongoing effort to produce post-quantum encryption standards. Whether they arrive in three years or 30, quantum computers loom on the horizon, and RSA, Diffie-Hellman, and other encryption schemes may be left vulnerable.

Finding a quantum-resistant cryptographic scheme isn't easy. Without a mathematical problem that is computationally hard, the last three decades of cybersecurity have played out like an increasingly intricate game, with researchers

perpetually building and breaking—or attempting to break—new candidates.

This push and pull has already emerged in the NIST post-quantum program. In February 2022, cryptographers found a fatal flaw in Rainbow, an algorithm that had survived three rounds of NIST's analysis. A few months later, after the NIST list had been winnowed again, Decru and his KU Leuven colleague Wouter Castryck announced that they'd broken another finalist, an algorithm called SIKE.

SIKE, which was developed a few years ago, was the brainchild of a collaboration among researchers and engineers at Amazon, Microsoft, the University of Versailles, and elsewhere. It is based on a special mathematical map, called an isogeny, that is made up of connections between elliptic curves. These maps can be turned into an encryption for communication, and outsiders can't eavesdrop without knowing the maps.

Related Story

What are quantum-resistant algorithms—and why do we need them?

When quantum computers become powerful enough, they could theoretically crack the encryption algorithms that keep us safe. The race is on to find new ones.

At Leuven, Decru and Castryck devise ways to use these so-called isogenies to build new, faster encryption approaches. They broke the most difficult version of SIKE in just a few hours of computing time using an ordinary desktop computer. (Since then, other groups have found ways to do it even faster.) What's more, Decru and Castryck did it almost accidentally, and only a few weeks after SIKE had been declared an alternate NIST finalist. "We weren't trying to break it at all," insists Decru. "We just tried to generalize it."

Chen says the case of SIKE—and Rainbow before it—illustrates a real-world tension that drives efforts to find quantum-proof algorithms. On one hand, she says, "you have to find a problem which is hard for both quantum computers and classical computers." On the other is implementation: transforming that hard problem into one that can be used in a real-world cryptographic system. Even with today's well-defined problems, Shorter says, it's very difficult to predict and prevent every loophole in every operating system and device on the market today. "And then there's interoperability testing and certifications and other tests," he says, "to make sure they are not only implemented correctly, but also securely."

The mathematical problem SIKE is based on seems computationally hard because there are so many different maps that could be constructed between curves. It may even be a one-way problem—and therefore quantum-proof. The flaw was in the design, which revealed too much of the transmitted information. Decru and Castryck cracked it because they inadvertently found a way to expose enough connecting points to give away the entire thing.

Other schemes have fared better. The first post-quantum encryption algorithm to be standardized, CRYSTALS-Kyber, delivers security through an approach that involves problems on lattices, mathematical objectsthat can be modeled as arrays of points. (There are five main families of post-quantum cryptographic methods. Isogeny and lattice approaches are two of them.)

CRYSTALS-Kyber is a general encryption scheme, like RSA, that can be used for tasks like securing online communication. Three other approved algorithms are designed to authenticate digital signatures, ensuring that digital documents haven't

been fraudulently signed. NIST plans to standardize these by spring 2024. Another three (it was four until SIKE was broken) could also be standardized in the next few years, as long as they survive further rounds of scrutiny.

But unless mathematicians can prove whether one-way functions exist, says Pass, the patterns that have always characterized cryptography will continue. “We’re back to this cat-and-mouse game, where it’s a game between algorithm designers proposing new candidate constructions and other designers trying to break them,” he says. Unless, of course, he—or someone in his field—can come up with an implementable, provably one-way function to settle the matter of encryption forever.

Until that time, cryptographers will remain in a messy limbo in which convincingly robust encryption schemes can be trusted—but only until they can’t.

The perfect math problem could take us out of this limbo, but it can’t be some sticky mess cooked up by an armchair algebraist over a long weekend. It must strike a balance between math and cryptography, with computational hardness on one side and easy implementation on the other. Stray too far from either of those properties, and it becomes vulnerable—if not now, then in the future. Hanging in the balance is the past, present, and future security of everyone’s data, everywhere. No pressure.

Stephen Ornes is a science writer based in Nashville.

by Stephen Ornes

The power of green computing

Sustainable computing practices have the power to both infuse operational efficiencies and greatly reduce energy consumption, says Jen Huffstetler, chief product sustainability officer at Intel.

By MIT Technology Review Insights

What's next for the world's fastest supercomputers

Scientists have begun running experiments on Frontier, the world's first official exascale machine, while facilities worldwide build other machines to join the ranks.

By Sophia Chen

AI-powered 6G networks will reshape digital interactions

The convergence of AI and communication technologies will create 6G networks that make hyperconnectivity and immersive experiences an everyday reality for consumers.

By MIT Technology Review Insights

How this Turing Award–winning researcher became a legendary academic advisor

Theoretical computer scientist Manuel Blum has guided generations of graduate students into fruitful careers in the field.

By Sheon Han

STAY CONNECTED

Illustration by Rose Wong