Bruce Schneier
Harvard University

NIST's Post-Quantum Cryptography Standards Competition

Quantum computing is a completely new paradigm for computers. A quantum computer uses quantum properties such as superposition, which allows a qubit (a quantum bit) to be neither 0 nor 1, but something much more complicated. In theory, such a computer can solve problems too complex for conventional computers.

Current quantum computers are still toy prototypes, and the engineering advances required to build a functionally useful quantum computer are somewhere between a few years away and impossible. Even so, we already know that that such a computer could potentially factor large numbers and compute discrete logs, and break the RSA and Diffie-Hellman public-key algorithms in all of the useful key sizes.

Cryptographers hate being rushed into things, which is why NIST began a competition to create a post-quantum cryptographic standard in 2016. The idea is to standardize on both a public-key encryption and digital signature algorithm that is resistant to quantum computing, well before anyone builds a useful quantum computer.

NIST is an old hand at this competitive process, having previously done this with symmetric algorithms (AES in 2001) and hash functions (SHA-3 in 2015). I participated in both of those competitions, and have likened them to demolition derbies. The idea is that participants put their algorithms into the ring, and then we all spend a few years beating on each other's submissions. Then, with input from the cryptographic community, NIST crowns a winner. It's a good process, mostly because NIST is both trusted and trustworthy.

In 2017, NIST received eighty-two post-quantum algorithm submissions from all over the world. Sixty-nine were considered complete enough to be Round 1 candidates. Twenty-six advanced to Round 2 in 2019, and seven (plus another eight alternates) were announced as Round 3 finalists in 2020. NIST was poised to make final algorithm selections in 2022, with a plan to have a draft standard available for public comment in 2023.

Cryptanalysis over the competition was brutal. Twenty-five of the Round 1 algorithms were attacked badly enough to remove them from the competition. Another eight were similarly attacked in Round 2. But here's the real surprise: there were newly published cryptanalysis results against at least four of the Round 3 finalists just months ago—moments before NIST was to make its final decision.

One of the most popular algorithms, Rainbow, was found to be completely broken. Not that it could theoretically be broken with a quantum computer, but that it can be broken today—with an off-the-shelf laptop in just over two days. Three other finalists, Kyber, Saber, and Dilithium, were weakened—with new techniques that will probably work against some of the other algorithms as well. (Fun fact: Those three algorithms were broken by the Center of Encryption and Information Security, part of the Israeli Defense Force. This represents the first time a national intelligence organization has published a cryptanalysis result in the open literature. And they had a lot of trouble publishing, as the authors wanted to remain anonymous.)

That was a close call, but it demonstrated that the process is working properly. Remember, this is a demolition derby. The goal is to surface these cryptanalytic results before standardization, which is exactly what happened. At this writing, NIST has chosen a single algorithm for general encryption and three digital-signature algorithms. It has not chosen a public-key encryption algorithm, and there are still four finalists. Check NIST's webpage on the project—https://csrc.nist.gov/projects/post-quantum-cryptography—for the latest information.

Ian Cassels, British mathematician and World War II cryptanalyst, once said that "cryptography is

continued on p. 107

Digital Object Identifier 10.1109/MSEC.2022.3184235
Date of current version: 6 September 2022

 continued from p. 108

a mixture of mathematics and muddle, and without the muddle the mathematics can be used against you." This mixture is particularly difficult to achieve with public-key algorithms, which rely on the mathematics for their security in a way that symmetric algorithms do not. We got lucky with RSA and related algorithms: their mathematics hinge on the problem of factoring, which turned out to be robustly difficult. Post-quantum algorithms rely on other mathematical disciplines and problems—code-based cryptography, hash-based cryptography, lattice-based cryptography, multivariate cryptography, and so on—whose mathematics are both more complicated and less well-understood. We're seeing these breaks because those core mathematical problems aren't nearly as well-studied as factoring is.

The moral is the need for cryptographic agility. It's not enough to implement a single standard; it's vital that our systems be able to easily swap in new algorithms when required. We've learned the hard way how algorithms can get so entrenched in systems that it can take many years to update them: in the transition from DES to AES, and the transition from MD4 and MD5 to SHA, SHA-1, and then SHA-3.

We need to do better. In the coming years we'll be facing a double uncertainty. The first is quantum computing. When and if quantum computing becomes a practical reality, we will learn a lot about its strengths and limitations. It took a couple of decades to fully understand von Neumann computer architecture; expect the same learning curve with quantum computing. Our current understanding of quantum computing architecture will change, and that could easily result in new cryptanalytic techniques.

The second uncertainly is in the algorithms themselves. As the new cryptanalytic results demonstrate, we're still learning a lot about how to turn hard mathematical problems into public-key cryptosystems. We have too much math and an inability to add more muddle, and that results in algorithms that are vulnerable to advances in mathematics. More cryptanalytic results are coming, and more algorithms are going to be broken.

We can't stop the development of quantum computing. Maybe the engineering challenges will turn out to be impossible, but it's not the way to bet. In the face of all that uncertainty, agility is the only way to maintain security. ■

Bruce Schneier is a security technologist, fellow, and lecturer at the Harvard Kennedy School, Cambridge, Massachusetts, USA, and the chief of security architecture of Inrupt, Inc. He blogs at www.schneier.com. Contact him at schneier@schneier.com.

Digital Object Identifier 10.1109/MSEC.2022.3199551