Quantum Computing and the Financial Services Industry

You may have heard about quantum computers and possibly read about the future threat quantum computers might pose to cryptography. What are quantum computers, and should the financial services industry as defined by ISO TC 68 be worried? These are the questions this document will try to address.

What's a Quantum Computer?

To fully understand what a quantum computer is probably requires a Ph.D. in quantum physics, but generally speaking, a quantum computer is a device that uses the quantum mechanical phenomena *superposition* and *entanglement* to solve certain specific problems *much* faster than classical computers.

Where a classical computer uses *bits*, which can be either 0 or 1, on or off, as its basic building block, a quantum computer uses *qubits* and a qubit carries much more information than a simple bit.[1] Qubits can exist in an intricate superposition between 0 and 1. This richness enables quantum computers, in principle, to solve certain specific problems dramatically faster than classical computers could. For example, a quantum computer could solve certain problems in days where classical computers would take years or more.

What Kind of Threat Would a Quantum Computer Pose to the Financial Services Industry?

While the promise of quantum computing is to solve certain specific problems dramatically faster than otherwise possible, currently quantum computers lack the stability and fault tolerance required to solve most real-world problems. As the technologies behind quantum computers mature, initially, quantum computers may provide advantages in AI and chemistry (see reference [9] and [10]), and in time, much more advanced quantum computers could potentially be a threat to the protections provided by today's cryptography.

Cryptography[2] is used pervasively throughout the financial services industry and it can be divided into asymmetric (e.g., public key) and symmetric key (e.g., T-DEA or AES) cryptography. Public key algorithms are used mainly for digital signatures and key establishment, and symmetric algorithms are used for encryption and message integrity. Examples of public key cryptography include the authentication and key establishment that takes place when establishing any secure Internet session, or the digital signatures that are applied to digital documents or for code signing. Examples of symmetric cryptography include the protection of PINs as they travel from PIN entry devices at ATMs or in stores,

[1] It takes two complex numbers to describe the state of one qubit.

[2] Defined as "a method for the transformation of data in order to hide its information content, prevent its undetected modification and/or prevent its unauthorized use" in [Source: ISO 7498-2]

through many hops to the card-issuing bank that will validate the PIN, or the encryption of data at rest in data centers across the industry and indeed the world.

A quantum computer's threat to symmetric cryptography would be much less serious than that to asymmetric cryptography.[3] Nonetheless, systems deploying symmetric cryptography would in general be advised in due time to migrate to AES (128 or 256) from T-DEA (see e.g. [7] and [8]). Although this by no means is a trivial task, the standards to support such a key migration are already in place [1], and the consequences for infrastructure is considered less impactful than that for public key algorithms[4]. Additional information may be found in [21].

In theory, a large scale, fault tolerant quantum computer leveraging Shor's algorithm [2] could break asymmetric cryptography based on either RSA or Elliptic Curve technology as used today. This is a serious threat since this technology forms the underpinning of practically all of today's public key cryptography.

However, and to the best of our knowledge, the current state of quantum computing is quite a number of steps away from realizing a quantum computer that can actually execute Shor's algorithm on sufficiently large input to break today's cryptography. How quickly this might change is a subject of continued debate and substantial disagreement amongst experts, with estimates ranging from 5 to 25 years or even more (see e.g. [16]).

What's the Expected Timeline for the Development of a Quantum Computer that Could Threaten Today's Cryptography?

The field of quantum computation is still not mature enough that we have a sound basis for predicting future developments. For the reader who wants to read 250+ pages thorough accounts of the state of quantum computing, [13] and [14] are both outstanding resources. The latter makes this statement: "***Key Finding 1****: Given the current state of quantum computing and recent rates of progress, it is highly unexpected that a quantum computer that can compromise RSA 2048 or comparable discrete logarithms-based public key cryptosystems will be built within the next decade*." The paper [16] by Mosca and Piani is also an excellent reference that provides the perspectives and predictions from a number of researchers.

Part of the problem of coming up with a firm timeline is that there are many different candidate approaches to building quantum computers and we don't yet know which approach will ultimately turn out to be best suited to create a quantum computer big enough to threaten today's cryptography.

[3] NIST for example in [6] states: "Taking these mitigating factors into account, it is quite likely that Grover's algorithm will provide little or no advantage in attacking AES, and AES 128 will remain secure for decades to come. Furthermore, even if quantum computers turn out to be much less expensive than anticipated, the known difficulty of parallelizing Grover's algorithm suggests that both AES 192 and AES 256 will still be safe for a very long time."

[4] This document does not discuss quantum key distribution (QKD), as QKD achieves much the same as a stream cipher, i.e. is in the domain of symmetric ciphers where the current state of affairs is largely satisfactory.

When you want to build a quantum computer, you have to decide on a way to implement the most basic components, the qubits (which represent information) and the quantum gates (which operate on the qubits). Among the different approaches, at the moment the most prominent ones are superconducting qubits, trapped ions, and photonics, each of which have their benefits and issues.

The development of a scalable quantum computer is still in the early stages with respect to the necessary research, design, engineering and manufacturing, and for example, the largest number of qubits anybody has managed to assemble into one computational architecture is 72 (Google's Bristlecone). The problem is that it only runs for milliseconds before failing which is not a sufficient period of time to solve real-world problems.

Why is it so hard to build a scalable quantum computer? The main reason is that the qubits in the quantum computer have to be kept nearly perfectly isolated from the surrounding environment. Any accidental interaction with the outside environment would change the qubit and destroy its state. At the same time, the qubits have to be connected to each other, in order to allow them to become entangled, and thereby achieve the quantum effect that causes the advantage of the quantum computer. Additionally, users have to be able to interact with them, in order to initialize them to a certain state, to control their operation from the outside and carry out the algorithm, and also to read out the result of their computation.

To reduce the impact of the environment on the qubits, most approaches to implementing qubits cool them down to temperatures very near absolute zero, requiring sophisticated dilution refrigerators. This in itself introduces new challenges in connecting the super cooled quantum part with the room temperature classical control software and hardware, which is used to initialize, control the operations, and read out the results, and avoiding that those connections disturb the qubits.

For these reasons, errors creep in everywhere in today's quantum computers and error correction schemes become necessary. The schemes contemplated today are of a kind that implement one (logical) qubit with thousands of physical qubits. Current (2019-2020) estimates, as exemplified in [17], are that a quantum computer with 20 million qubits and 2.7 billion Toffoli gates is needed to break a 2048 bit RSA modulus in 8 hours.

We are thus looking at going from ~100 physical qubits to at least tens of millions, before any threat to our current asymmetric cryptography will materialize. However, the area is certainly dynamic and sufficiently uncertain that it will be prudent to pay close attention to its development.

Professor Mosca from the University of Waterloo in Canada has developed a "theorem" that can be quickly tailored to any type of information or data that needs to be protected. The theorem is:

If we continue to use our current cryptography for time X and the data needs to be secure for time Y then we are in trouble if X+Y>Z, the time till a quantum computer exists that can break the cryptography.

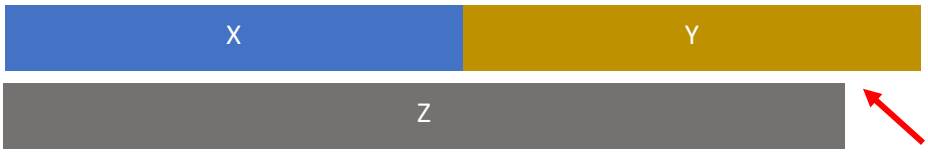

For some early warning signs to look out for, see the Appendix.

What are the defensive measures preparing for the emergence of scalable quantum computing?

Based on their concern about the emergence of scalable, fault tolerant quantum computers, NIST in 2016 started an initiative [15] to solicit so-called post-quantum algorithms for eventual standardization. They solicited public key algorithms in the areas of digital signature, encryption, and key-establishment. By the deadline of November 30, 2017 NIST had received over 80 submissions.

The period from December 2017 to now has been used for peer review of the proposals, followed by rounds of selection by NIST. Through that process, there are now 7 finalists and 8 alternates across two areas, (i) Digital Signatures and (ii) Public Key Encryption and Key Establishment. The final publication of the NIST candidates is scheduled for December 2021. NIST expects to release their draft standards for post-quantum algorithms in the 2022-2024 timeframe, after which time ISO can be expected to finalize its standardization. For an introduction to the different post-quantum systems, refer to ISO/IEC JTC1 SC27 WG2's standing document SD8 ([19]).

Conclusion

Given the future disruptive potential of quantum computing, now would be a good time for financial organizations to start assessing their exposure to algorithms put at risk by quantum computers and start thinking about how a migration to post-quantum algorithms can proceed. They can also start evaluating the post-quantum algorithms being finalized by NIST and begin to analyze the impact each would have on their environment.

From a practical implementation perspective, it can be noted that each of the current NIST finalists involves key and cryptogram objects that are larger than currently in use in financial services and there may be additional performance or resource impacts. When NIST standards emerge, followed by ISO standards and guidance specifically for the financial services industry, migration plans can solidify and be ready for the time when any migration would actually need to occur.

Quoting from *Quantum Computing: Progress and Prospects* [14] "*Even if a quantum computer that can decrypt current cryptographic ciphers is more than a decade off, the hazard of such a machine is high enough—and the time frame for transitioning to a new security protocol is sufficiently long and uncertain—that prioritization of the development, standardization, and deployment of post-quantum cryptography is critical for minimizing the chance of a potential security and privacy disaster*".

Cryptographic migrations in the financial services industry take a long time, but comprehensive and early planning and preparation can make the effort much more manageable.

References

1. ISO 9564, Financial Services – Personal Identification Number (PIN) management and security.
2. P.W. Shor, Polynomial-time algorithms for prime factorization and discrete logatithms on a quantum computer, SIAM J. Sci. Statist. Comput. 26 (1997) 1484. https://arxiv.org/abs/quant-ph/9508027 .
3. Evan Jeffrey, The physics of building a quantum computer, Real World Crypto 2017, https://www.youtube.com/watch?v=sQMm4w23mhw .

4. A.G. Fowler, M/ Mariantoni, J.M. Martinis, A.N. Cleland, Surface Codes: Towards practical large-scale quantum computation, DOI:10.1103/PhysRevA.86.032324 , https://arxiv.org/ftp/arxiv/papers/1208/1208.0928.pdf , 2012.

5. O. Nissim, et al.: Demonstrating Quantum Error Correction that Extends the Lifetime of Quantum Information, Feb 2016, https://arxiv.org/pdf/1602.04768.pdf.

6. NIST FAQ, https://csrc.nist.gov/projects/post-quantum-cryptography/faqs .

7. PCI SCC, PCI PIN Security, Requirements and Testing Procedures, Version 3.0, August 2018, https://www.pcisecuritystandards.org/documents/PCI_PIN_Security_Requirements_Testing_v3_Aug2018.pdf .

8. ISO/TR 14742 Financial Services – Recommendations on cryptographic algorithms and their use.

9. K. Bourzac, Chemistry is quantum computing's killer app, Chemical and Engineering News, Vol. 95 Issue 43, pp. 27-31, Oct. 30, 2017, https://cen.acs.org/articles/95/i43/Chemistry-quantum-computings-killer-app.html .

10. M. Schuld, F. Petruccione, *Supervised Learning with Quantum Computers*, Springer Verlag 2018.

11. Arute, F. et al., Quantum supremacy using a programmable superconducting processor, Nature **574**, pages 505–510(2019), https://www.nature.com/articles/s41586-019-1666-5 .

12. Yirka, B. Chinese photonic quantum computer demonstrates quantum supremacy, Phys.org, Dec. 4, 2020: https://phys.org/news/2020-12-chinese-photonic-quantum-supremacy.html .

13. BSI, German Federal Office for Information Security, Status of quantum computer development, version 1.2, https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Publikationen/Studien/Quantencomputer/P283_QC_Studie-V_1_2.html;jsessionid=B4CA190A1267FFC62EB36F9A048B58C6.1_cid500

14. E. Grumbling, M. Horowitz, eds. *Quantum Computing: Progress and Prospects*, from The National Academy of Sciences, The National Academic Press, 2019. Available in preprint from https://www.nap.edu/catalog/25196/quantum-computing-progress-and-prospects.

15. National Institute for Standards and Technology, Post-Quantum Initiative, https://csrc.nist.gov/projects/post-quantum-cryptography .

16. M. Mosca, M. Piani, *Quantum threat timeline*, The Global Risk Institute, 2020, https://globalriskinstitute.org/publications/quantum-threat-timeline-report-2020/ .

17. C. Gidney, M. Ekerå, *How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits*, https://arxiv.org/abs/1905.09749 .

18. T. Häner, S. Jaques, M. Naehrig, M. Roetteler, M. Soeken, *Improved quantum circuits for elliptic curve discrete logarithms*, https://arxiv.org/abs/2001.09580 .

19. ISO-IEC/JTC 1 SC27/WG2, Standing Document 8, Post Quantum Cryptography, Parts 1-6, https://www.din.de/resource/blob/721042/4f1941ac1de9685115cf53bc1a14ac61/sc27wg2-sd8-data.zip or https://www.din.de/en/meta/jtc1sc27/downloads .

20. V. Gheorghiu, M. Mosca, A resource estimation framework for quantum attacks against cryptographic functions – recent developments, GRI quantum risk assessment report, Feb. 2020, https://globalriskinstitute.org/download/gri-quantum-risk-assessment-report-part-5-report/ .

21. ISO/IEC DIS 18033-1 Information security – encryption algorithms – Part 1: General, Under development.

Appendix - What's the Canary in the Coal Mine?

It will be worth paying attention to factorization records (e.g. of 50-digit numbers or larger), or even the first announcement of an error correcting or fault tolerant quantum computer with dozens of *logical* qubits. The absence of such announcements would indicate that we still have a considerable time before current day public key cryptography is threatened. Lately there have been a couple of announcements of so-called quantum supremacy [11], [12], but the problems solved in those efforts are specifically designed with this purpose in mind and do not easily translate into breaking cryptography.

We have seen significant optimizations in implementations of Shor's algorithm, with [17] and [18] as examples. The optimizations so far have meant an order of magnitude smaller quantum computer is needed (tens of millions of physical qubits, not hundreds of millions) to break current RSA or elliptic curve cryptography, but rigorous error correction is still required. If we were to see more fundamental changes to Shor's algorithm that made it tolerate more errors, that would lower the requirements to a scalable quantum computer and bring it closer to reality.

One very significant signal would be the realization of stable quantum storage – even just one stable logical qubit would most likely require thousands of physical qubits with error correction but would be an important foundational building block towards a scalable quantum computer.

One other issue to consider is that besides the public efforts discussed so far, several countries may also have national quantum computing efforts which are not public. We would not expect any public signals to come out of such efforts.