Christian Schaffner
QuSoft and ILLC
University of Amsterdam
c.schaffner@uva.nl

Recent developments in quantum cryptography

Christian Schaffner is a NWO Vidi Laureate and assistant professor at the ILLC at UvA, where he works on quantum cryptography. Quantum cryptography is the art and science of exploiting quantum mechanical effects in order to perform cryptographic tasks. Recent progress in building quantum computers has led to new opportunities for cryptography, but also endangers existing cryptographic schemes. In this article Schaffner surveys both aspects of this double-edged sword.

Quantum supremacy

We stand at a special moment in time, where the promise of a quantum computer and quantum internet seems to be within reach. Large-scale efforts around the world by academic and industrial players including many heavyweights such as Google, IBM, Microsoft and Intel are in progress to develop quantum technologies. The first experiments demonstrating 'quantum supremacy' are expected to be carried out possibly as early as at the end of 2017. Such experiments will demonstrate that a quantum computer can solve some (probably uninteresting) problem faster than any classical supercomputer. In order to outperform today's biggest classical computers, a quantum computer has to work with at least 40 to 50 *quantum bits* (also called *qubits*). In contrast to a classical bit which can take on values either 0 or 1, a quantum bit can be in an arbitrary superposition between 0 and 1. Mathematically, the (pure) state of a single qubit can be expressed as unit vector in a two-dimensional complex Hilbert space $\mathcal{H}$, i.e.

$$|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$$
$$\text{with } \alpha,\beta \in \mathbb{C} \text{ and } |\alpha|^2 + |\beta|^2 = 1$$

where the two states $|0\rangle, |1\rangle$ form an orthonormal basis of $\mathcal{H}$, and α and β are called amplitudes of the state. Measuring the quantum state $|\phi\rangle$ in basis $|0\rangle, |1\rangle$ yields a classical bit b as outcome. Outcome $b=0$ is obtained with probability $|\alpha|^2$ and outcome $b=1$ is obtained with probability $|\beta|^2$.

Quantum mechanics dictates that the state space of n qubits is given by the n-fold tensor product $\mathcal{H}^{\otimes n}$. Remarkably, this space has dimension 2^n which is exponential in the number of qubits! Therefore, the state of a 40-qubit quantum computer can be described by 2^{40} complex amplitudes, and a quantum computation on such a device can be classically simulated by tracking the changes of all these amplitudes during the computation. The exponential size of the state space explains why it becomes intractable to classically simulate arbitrary quantum computations on, say, more than 50 qubits.

One of the most promising proposals by the Google group [4] for an experiment which demonstrates the superiority of a quantum computer over all classical computers is to perform the following quantum computation: initialize a two-dimensional grid of 7 by 7 qubits in a fixed state, and apply a number of random 2-qubit quantum gates on any two neighboring qubits. The squared amplitudes of the quantum state after the application of the quantum circuit define a probability distribution, and repeated application of the same quantum procedure allows to sample from this distribution by simply measuring the final quantum state. There is complexity-theoretic evidence that it is difficult to sample from the same distribution on a classical computer [1, 4].

Clearly, such a sampling problem is not of great practical interest. However, it is well-known since the late eighties that large-scale quantum computers can solve some important problems far more efficiently than classical computers. The most well-known example is the problem of finding the prime factorisation of an integer $N \in \mathbb{N}$. For large problem sizes, the fastest classical algorithms have an expected running time $O(e^{(\log N)^{1/3}})$ which is subexponential in the size of the integer N to be factored. In contrast, a quantum algorithm discovered by Peter Shor requires only $O((\log N)^2)$ quantum operations to solve the problem. This time complexity is polynomial in the size of N and therefore exponentially faster than the best-known classical algorithm.

The main high-level idea of Shor's quantum algorithm is to reduce the factoring problem to the problem of period-finding, which can be solved efficiently on a quantum computer by exploiting the power of the quantum Fourier transform. It turns out that also other cryptographically relevant problems such as the discrete-logarithm problem in finite fields can be reduced to period-finding, and hence fall to Shor's quantum algorithm as well, see the box on the next page.

Discrete logarithms

For a finite group G of order q, generated by an element $g \in G$, the *discrete-logarithm problem* is defined as follows: Given a uniformly random group element $h \in G$, find x such that $g^x = h$. For a generic group G (i.e. if one can only use the group operation, and no additional mathematical structure of G), one can prove that any classical algorithm must perform $\Omega(\sqrt{q})$ group operations to compute discrete logarithms [18]. Again, Shor's algorithm can solve the problem even in a generic group using a number of quantum operations which is polynomial in $\log(q)$. Popular choices of the finite group G in cryptography are $\mathbb{Z}_p^*$, the multiplicative group of integers modulo a large prime p (with some additional requirements), or the elliptic-curve groups over finite fields. The advantage of elliptic curves over $\mathbb{Z}_p^*$ is that instance sizes can be much smaller while keeping the conjectured difficulty of the discrete-logarithm problem at the same level.

Quantum-safe cryptography

Cryptography is of utmost importance in today's digital world. Cryptographic systems that use classical communication are widely deployed and form the cornerstone of digital technologies, ranging from the internet to mobile phones. If present-day cryptographic systems were broken, we could no longer securely perform online banking, use mobile devices or the internet of things, our medical data would no longer be protected, et cetera. Almost all of the currently employed public-key cryptography is based on either the factoring or on the discrete-logarithm problem. Therefore, these systems can all be broken using Shor's quantum algorithm. An example is the widely deployed RSA cryptosystem [15]. Forging RSA signatures on operating-system updates would allow an attacker to gain control of billions of computers around the world.

Crucially, the security of such cryptographic systems can be broken *retroactively* by an attacker who obtains a large-scale quantum computer in the future. Data encrypted today can thus lose its confidentiality once a quantum computer becomes available. For the protection of state secrets or medical data, new forms of quantum-safe cryptographic protection have to be employed *already today* in order to guarantee the confidentiality of the data for the coming decades.

At the moment, large parts of the international cryptographic research community are searching for new public-key schemes based on mathematical problems which are believed to be hard also for quantum computers. The US National Institute for Standards in Technology (NIST) has initiated a project to determine a new standard for such quantum-safe scheme, see [14]. (The term 'post-quantum cryptography' is quite well-established for these quantum-safe schemes, but chosen somewhat unfortunately, because the research area is concerned with cryptography which is still secure *at the beginning* and not after the end of the era of large-scale quantum computers.)

The main contenders in terms of assumptions for classical quantum-safe cryptography are the following:

- lattice-based cryptography (see article by Léo Ducas in this issue),
- multivariate cryptography,
- hash-based cryptography,
- code-based cryptography,
- supersingular elliptic curve isogeny cryptography (see article by Mathijs Coster in this issue).

In order to evaluate the proposed schemes and assumptions, an extensive amount of quantum cryptoanalysis will be performed over the coming years. The challenge is that expertise from very different research fields such as cryptography, mathematics and quantum computing is required. The Quantum Software Consortium, a collaboration of researchers from CWI, QuSoft, TU Delft and Leiden university have recently obtained a large gravitation grant from NWO (18.8 million Euro for ten years). As this consortium unites experts from all aforementioned fields, it is natural that a considerable part of the grant will be spent on the quantum cryptanalysis of candidate schemes for quantum-safe cryptography. A thorough understanding of the capabilities of current quantum algorithms and quantum computers is essential to determine the correct parameters for which the above hardness assumptions hold.

It is important to note that the use of quantum technology cannot only be used to attack existing classical schemes, but on the contrary, it has been realized already in the late sixties by Stephen Wiesner [19] that the use of quantum communication can lead to new cryptographic applications. The most famous example is Quantum Key Distribution (QKD) [3, 12] which allows two honest parties Alice and Bob to establish a classical secret key even in the presence of an eavesdropper Eve who has complete control over the quantum communication between Alice and Bob, and who can eavesdrop on (but not modify) the classical communication. It can be proven in an information-theoretic sense that no matter how much computational power an eavesdropper has, her knowledge about the secret key obtained by Alice and Bob can be made arbitrarily small. It is quite remarkable that such an information-theoretical key establishment is possible using quantum communication, because for purely classical communication, it is excluded by Shannon's impossibility result for perfectly secure encryption [16].

On a higher level, quantum key distribution is yet another form of quantum-safe cryptography which can be compared to the other proposals above. The main advantage of QKD is the information-theoretic security of the established keys, and the fact that an attacker cannot retroactively break the security of QKD. An attacker has to act at the moment when the key-establishment protocol is performed, so large-scale quantum computers developed in the future do not change the security of keys established today. The main disadvantages of QKD are the high deployment costs (because quantum communication is required) and the current distance limit of about 150 kilometers between Alice and Bob (which implies that for larger distances, a multi-hop connection is required where the middle nodes have to be trusted).

Quantum fully-homomorphic encryption

In recent research [11], we were able to develop a new method for securely delegating a quantum computation. In fact, we developed the quantum analogue of what is classically known as fully-homomorphic encryption. Homomorphic encryption is a special form of encryption and is best motivated by the scenario of secure cloud computing, where a user Alice would like to outsource a difficult or tedious computation to an untrusted computing cloud.

For instance, Alice would like to tag a lot of holiday pictures, but instead of doing this tedious job manually, she would like to delegate the task to an advanced image-recognition algorithm in the cloud. As Alice does not trust the cloud with all her private images, she can use homomorphic encryption to encrypt her pictures and send them to the cloud, which cannot see the contents due to the secrecy of the encryption but which can nevertheless run its tagging algorithms on the encrypted images due to the homomorphic property of the encryption. As a result, the cloud returns encrypted tags of Alice's pictures, and Alice is the only person who can decrypt them. For a long time, cryptographers thought it was impossible to construct such homomorphic encryption systems, until Craig Gentry showed in a seminal article in 2009 how to do this [13]. This breakthrough sparked an enormous amount of cryptographic research on the topic and led in recent years to another hot topic of cryptographic research on (software) obfuscation.

Given the recent progress building a quantum computer, it is very natural to consider the scenario of a quantum-computing cloud. In fact, it is rather likely that quantum computing will only be available at a few quantum-computing facilities in the world. As a concrete example, IBM is already providing access to a mini quantum-computation cloud by allowing the general public to run algorithms on their 5-qubit computer. Hence, we were wondering whether it is possible to construct homomorphic encryption schemes for quantum data. Previous results in this direction had serious drawbacks and allowed for homomorphic evaluations of only a very limited set of operations. Our recent result [11] provides quantum encryptions that allow homomorphic evaluations of arbitrary efficient quantum circuits.

Illustration: Tremani, TU Delft

Figure 1 In early stages of the quantum age, quantum computing will only be available in specialized facilities. Secure delegated quantum computation allows to outsource a quantum computation to a server in the quantum-computing cloud without revealing the inputs.

Our quantum fully-homomorphic encryption scheme offers a round-optimal solution to the problem of *secure delegated quantum computation*, see Figure 1. Pioneering work of Childs [10] and Arrighi and Salvail [2] studied this problem for the first time. The first practical and universal protocol for private delegated quantum computation, called 'universal blind quantum computation', was given by Broadbent, Fitzsimons and Kashefi [5]. In this protocol, the client only needs to be able to prepare random single-qubit auxiliary states, but does not require quantum memory nor a quantum processor. Via a classical interaction phase, the client remotely drives a quantum computation of her choice, such that the quantum server cannot learn any information about the computation that is performed — with only the client learning the output. Our new scheme from [11] reduces the number of rounds of interaction between client and server, at the expense of some extra demands on the quantum capabilities of the client.

Currently, we are working on extending our scheme so that the client can classically verify whether the server has carried out the correct quantum computation. Such a property seems crucial for using it as a building block to obtain more advanced cryptographic primitives such as zero-knowledge proofs [8], quantum one-time programs [6], or quantum obfuscation.

Conclusion

These are exciting times for the active research field of quantum cryptography, as quantum computers are becoming a reality in the near future. Many challenges await us in finding alternatives to the currently broken classical public-key cryptosystems, and in further exploiting the power of quantum computing to extend cryptographic schemes into the quantum domain.

Acknowledgments

Small parts of this article have appeared previously in a survey article co-authored by Anne Broadbent [9], where we introduce quantum cryptography to classical cryptographers. I would like to thank Anne Broadbent for her kind permission to reuse some parts here. I am also grateful for comments on this article from Anne Broadbent and Ronald de Wolf.

References

1 S. Aaronson and L. Chen, Complexity-theoretic foundations of quantum supremacy experiments, arxiv:1612.05903.

2 P. Arrighi and L. Salvail, Blind quantum computation. *Int. J. Quantum Inf.* 4 (2006), 883–898.

3 C.H. Bennett and G. Brassard, Quantum cryptography: Public key distribution and coin tossing, *International Conference on Computers, Systems and Signal Processing*, 1984, pp. 175–179.

4 S. Boixo, S.V. Isakov, V.N. Smelyanskiy, et al., Characterizing quantum supremacy in near-term devices, arxiv:1608.00263.

5 A. Broadbent, J. Fitzsimons and E. Kashe, Universal blind quantum computation, in *FOCS 2009*, pp. 517–526.

6 A. Broadbent, G. Gutoski and D. Stebila, Quantum one-time programs, in *CRYPTO 2013*, pp. 344–360.

7 A. Broadbent and S. Jeery, Quantum homomorphic encryption for circuits of low T-gate complexity, in *CRYPTO 2015*, pp. 609–629.

8 A. Broadbent, Z. Ji, F. Song and J. Watrous, Zero-knowledge proof systems for QMA, in *FOCS 2016*, IEEE, 2016, pp. 31–40.

9 A. Broadbent and C. Schaner, Quantum cryptography beyond quantum key distribution, *Designs, Codes and Cryptography* 78 (2016), 351–382.

10 A. Childs, Secure assisted quantum computation, *Quantum Information and Computation* 5 (2005), 456–466.

11 Y. Dulek, C. Schaner and F. Speelman, Quantum homomorphic encryption for polynomial-sized circuits, in *CRYPTO 2016*, pp. 3–32.

12 A.K. Ekert, Quantum cryptography based on bell's theorem, *Physical Review Letters* 67 (1991), 661–663.

13 C. Gentry, Fully homomorphic encryption using ideal lattices, in *STOC 2009*, Vol. 9, pp. 169–178.

14 National Institute of Standards and Technology, Post-Quantum crypto Project, http://csrc.nist.gov/groups/ST/post-quantum-crypto.

15 R.L. Rivest, A. Shamir and L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, *Communications of the ACM* 21(2) (1978), 120–126.

16 C.E. Shannon, Communication theory of secrecy systems, *The Bell System Technical Journal* 28(4) (1949), 656–715.

17 P.W. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *35th Annual Symposium on Foundations of Computer Science – FOCS 1994*, IEEE, 1994, pp. 124–134.

18 Victor Shoup, Lower bounds for discrete logarithms and related problems, in *EUROCRYPT 1997*.

19 S. Wiesner, Conjugate coding. *SIGACT News* 15(1) (1983), 78-88, originally written in 1968 but unpublished.