Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 4, 2017

P. Hoffman
ICANN

May 3, 2017

The Transition from Classical to Post-Quantum Cryptography
draft-hoffman-c2pq-00

Abstract

Quantum computing is the study of computers that use quantum features in calculations. For over 20 years, it has been known that if large-scale quantum computers could be built, they could have a devastating effect on classical cryptographic algorithms such as RSA and elliptic curve signatures and key exchange, as well as on encryption algorithms. There has already been a great deal of study on how to create algorithms that will resist large-scale quantum computers, but so far, the properties of those algorithms make them onerous to adopt before they are needed.

Small-scale quantum computers are being built today, but it is still far from clear when large-scale quantum computers that can be used to break classical algorithms with key sizes commonly used today will be available. It is important to be able to predict when large-scale quantum computers usable for cryptanalysis will be possible so that organization can change to post-quantum cryptographic algorithms well before they are needed.

This document describes quantum computing, how it can be used to attack classical cryptographic algorithms, and possibly how to predict when large-scale quantum computers will become feasible.

This Internet-Draft will expire on November 4, 2017.

Table of Contents

1. Introduction

Early drafts of this document use "@@@@@" to indicate where the authors particularly want input from reviewers. The authors welcome all types of review, but the areas marked with "@@@@@" are in the most noticeable need of new material. (The authors particularly appreciate new material that comes with references that can be included in this document as well.)

1.1. Disclaimer

**** This is the very first version of this draft. **** As such, it has had absolutely no review in the cryptography community. Statements in this document might be wrong; given that the entire document is about cryptography, those wrong statements might have significant security problems associated with them.

Readers of this document should not rely on any statements in this version of this draft. As the draft gets more input from the cryptography community over time, this disclaimer will be softened and eventually eliminated.

1.2. Executive Summary

The development of quantum computers that can break classical cryptographic keys is at a very early stage. None of the published examples of such quantum computers is useful in breaking keys that are in use today. There is a great amount of interest in this development, and researchers expect large strides in this development in the coming decade.

Because the world does not know when large-scale quantum computers that can break cryptographic keys will be available, organizations should be watching this so that they have plenty of time to either change to larger key sizes for classical cryptography or to change to post-quantum algorithms. See Section 5 for a fuller discussion of determining how to predict when large-scale quantum computers might become feasible.

1.3. Terminology

The term "classical cryptography" is used to indicate the cryptographic algorithms that are in common use today. In particular, signature and key exchange algorithms that are based on the difficulty of factoring numbers into two large prime numbers, or

are based on the difficulty of determining the discrete log of a large composite number, are considered classical cryptography.

The term "post-quantum cryptography" is the invention and study of signature and key exchange algorithms that are not based on the difficulty of factoring numbers into two large prime numbers, nor on the difficulty of determining the discrete log of a large composite number.

Note that these definitions apply to only one aspect of quantum computing as it relates to cryptography. It is expected that quantum computing will also be able to be used agains symmetric key cryptography to make it possible to search for a secret symmetric key using far fewer operations than are needed using classical computers (see Section 4 for more detail). However, using longer keys to thwart that possibility is not normally called "post-quantum cryptography".

There are many terms that are only used in the field of quantum computing, such as "qubit", "quantum algorithm", and so on. Chapter 1 of [NielsenChuang] has good definitions of such terms.

The "^" symbol is used to indicate "the power of". The term "log" always means "logarithm base 2".

1.4. Not Covered: Post-Quantum Cryptographic Algorithms

This document discusses when an organization would want to consider using post-quantum cryptographic algorithms, but definitely does not delve into which of those algorithms would be best to use. Post-quantum cryptography is an active field of research; in fact, it is much more active than the study of when we might want to transition from classical to post-quantum cryptography.

Readers interested in post-quantum cryptographic algorithms will have no problem finding many articles proposing such algorithms, comparing the many current proposals, and so on. An excellent starting point is the web site <http://pqcrypto.org/>. Another is the article on post-quantum cryptography at Wikipedia: <https://en.wikipedia.org/wiki/Post-quantum_cryptography>.

In addition, various organizations are working on standardizing the algorithms for post-quantum cryptography. For example, the US National Institute of Standards and Technology (commonly just called "NIST") is holding a competition to evaluate post-quantum cryptographic algorithms. NIST's description of that effort is currently at <http://csrc.nist.gov/groups/ST/post-quantum-crypto/>.

1.5. Not Covered: Quantum Cryptography

Outside of this section, this document does not cover "quantum cryptography". The field of quantum cryptography is related to quantum computers, but not to cryptanalysis. Quantum cryptography is used to share random values that cannot be observed by outside parties without discovery.

1.6. Where to Read More

There are many reasonably accessible articles on Wikipedia, notably <https://en.wikipedia.org/wiki/Quantum_computing>.

@@@@@ Note to the CFRG: please review the various pages at Wikipedia and update them if they are wrong or out of date. Doing so is incredibly helpful to the world.

[NielsenChuang] is a well-regarded college textbook on quantum computers. Prerequisites for understanding the book include linear algebra and some quantum physics; however, even without those, a reader can probably get value from the introductory material in the book.

@@@@@ Maybe add more references that might be useful to non-experts.

2. Brief Introduction to Quantum Computers

A quantum computer is a computer that uses quantum bits (qubits) in quantum circuits to perform calculations. Quantum computers also use classical bits and regular circuits: most calculations in a quantum computer are a mix of classical and quantum bits and circuits.

@@@@@ This can be expanded and made less hand-wavy.

Qubits are valuable in quantum computers when they are combined in calculations. Combining qubits in a calculation requires that the qubits are correlated. Correlating qubits requires much more effort than correlating classical bits (such as in registers or volatile memory), which is one of the main reasons that developing quantum computers has proven more difficult than early development of classical computers.

@@@@@ Discuss measurements and how they have to be done with correlated qubits.

2.1. Quantum Computers that Discover Cryptographic Keys

Quantum computers are expected to be useful in many applications in the future. However, this document only discusses how they might be used to discover cryptographic keys faster than classical computers. In order to discover cryptographic keys, a quantum computer needs to have a quantum circuit specifically designed for the type of key it is attempting to break.

In order for a quantum computer to be useful to discover the type and size keys that are in common use today, it has to have a circuit with thousands of qubits. Smaller quantum computers (those with fewer qubits and simpler circuits) don't speed up cryptanalysis of these keys at all. That is, no one has devised a way to combine a bunch of smaller quantum computers to perform the same attacks on cryptographic keys as a properly-sized quantum computer.

This is why this document uses the term "large-scale quantum computer" when describing ones that can be used to break keys: there will certainly be small-scale quantum computers built first, but those computers cannot be used to discover the type and size keys that are in common use today.

2.2. Qubits, Error Detection, and Error Correction

@@@@@ Lots of material goes here. We will need recent references for how many physical qubits are needed for each corrected qubit. It's OK if this section has lots of references, but hopefully they don't contradict each other.

2.3. Physical Designs for Quantum Computers

Quantum computers can be built using many different physical technologies. Deciding which physical technologies are best to pursue is an extremely active research topic. A few physical technologies (particularly trapped ions, super-conduction using Josephson junctions, and nuclear magnetic resonance) are currently getting the most press, but other technologies are also showing promise.

@@@@@ It would be useful to have maybe two paragraphs about each physical design that is being actively pursued.

2.4. Challenges for Physical Designs

Different designs have different challenges to overcome before the physical technology can be scaled enough to build a useful large-scale quantum computer. Some of those challenges include the

following. (Note that some items on this list apply only to some of the physical technologies

Temperature: Getting stable operation without extreme cooling is difficult for many of the proposed technologies. The definition of "extreme" is different for different low-temperature technologies.

Stabilization: The length of time every qubit in a circuit holds is value

Quantum control: Coherence and reproducibility of qubits

Error detection and correction: Getting accurate results through simultaneous detection of bit-flip and phase-flip. See Section 2.2 for a longer description of this.

Substrate: The material on which the qubit circuits are built. This has a large effect on the stability of the qubits.

Particles: The atoms or sub-atomic particles used to make the qubits

Scalability: The ability to handle the number of physical qubits needed for the desired the circuit

Architecture: Ability to change quantum gates in a circuit

3. Quantum Computers and Public Key Cryptography

The area of quantum computing that has generated the most interest in the cryptographic community is the ability of quantum computers to find the secret keys in the RSA and Diffie-Hellman algorithms using many fewer operations than classical computers would need to use. It is widely believed that factoring large numbers and finding discrete logs using classical computers increases with the exponential size of the key. [RFC3766] describes in detail how classical computers can be used to determine keys; even though that RFC is over a decade old, no significant changes have been made to the process of classical attacks on RSA and Diffie-Hellman. @@@@@ CFRG: is that true? Does RFC 3766 need to be updated?

Shor's algorithm shows that these problems can be solved on quantum computers in polynomial time, meaning that the speed of finding the keys is a polynomial function based on the size of the keys, which would require significantly fewer steps than a classical computer. The definitive paper on Shor's algorithm is [Shor97].

3.1. Explanation of Shor's Algorithm

@@@@@ Pointers to understandable articles would be good here.

@@@@@ Describe period-finding and why it applies to finding prime factors and discrete logs.

@@@@@ Give the steps for applying Shor's algorithm to 2048-bit RSA. Describe how many rounds of the quantum subroutine would likely be needed. Describe how many rounds of the classical loop would likely be needed.

@@@@@ Give the steps for applying Shor's algorithm to 256-bit elliptic curves. Describe how many rounds of the quantum subroutine would likely be needed. Describe how many rounds of the classical loop would likely be needed.

3.2. Properties of Large-Scale Quantum Computers Needed for Discovering Public Keys

Researchers have built small-scale quantum computers that implement Shor's algorithm, factoring numbers with four or five bits. These are used to show that Shor's algorithm is possible to realize in actual hardware.

@@@@@ References are needed here. Did they implement all of Shor's algorithm, including the looping logic in the classical part and the looping logic in the quantum part?

@@@@@ Numbers and explanation is needed below:

A quantum computer that can determine the secret keys for 2048-bit RSA would require SOME NUMBER GOES HERE correlated qubits and SOME NUMBER GOES HERE circuit elements. A quantum computer that can determine the secret keys for 256-bt elliptic curves would require SOME NUMBER GOES HERE correlated qubits and SOME NUMBER GOES HERE circuit elements.

4. Quantum Computers and Symmetric Key Cryptography

Section 3 is about Shor's algorithm and compromises to public key cryptography. There is a second quantum computing algorithm, Grover's algorithm, that is often mentioned at the same time as Shor's algorithm but, with respect to cryptanalysis, only applies to symmetric ciphers such as AES. The definitive paper on Grover's algorithm is by Grover: [Grover96]. Grover later wrote a more accessible paper about the algorithm in [QuantumSearch].

Grover's algorithm gives a way to search for keys to symmetric algorithms in the square root of the time that a normal exhaustive search would take. Thus, a large-scale quantum computer that implemented Grover's algorithm could find a secret AES-128 key in about 2^64 steps instead of the 2^128 steps that would be required for a classical computer.

When it appears that it is feasible to build a large-scale quantum computer that can defeat a particular symmetric algorithm at a particular key size, the proper response would be to use keys with twice as many bits. That is, if one is using the AES-128 algorithm and there is a concern that an adversary might be able to build a large-scale quantum computer that is designed to attack AES-128 keys, move to an algorithm that has keys twice as long as AES-128, namely AES-256.

It is currently expected that large-scale quantum computers that implement Grover's algorithm are expected to be built long before ones that implement Shor's algorithm are. There are two primary reasons for this:

- Grover's algorithm is likely to be useful in areas other than cryptography. For example, a large-scale quantum computer that implements Grover's algorithm might be used to help create medicines by speeding up complex problems that involve how proteins fold. @@@@@ Add more likely examples and references here.

- A large-scale quantum computer that can be used to break AES-128 will likely much smaller (and thus easier to build) than one that implements Shor's algorithm for 256-bit elliptic curves or 2048-bit RSA/DSA keys.

4.1. Explanation of Grover's Algorithm

@@@@@ Give the steps for applying Grover's algorithm to AES-128.

4.2. Properties of Large-Scale Quantum Computers Needed for Discovering Symmetric Keys

@@@@@ Numbers and explanation is needed below:

A quantum computer that can determine the secret keys for AES-128 would require SOME NUMBER GOES HERE correlated qubits and SOME NUMBER GOES HERE circuit elements.

@@@@@ <https://arxiv.org/abs/1512.04965> indicates that the quantum part of the computer would have more than 2^80 quantum gates, which might be prohibitive for physical hardware.

5. Predicting When Useful Cryptographic Attacks Will Be Feasible

If quantum computers that perform useful cryptographic attacks can be built in the future, many organizations will want to start using post-quantum algorithms well before those computers can be built. However, given how few implementations of such quantum computers exist (even for tiny keys), it is impossible to predict with any accuracy when quantum computers that perform useful cryptographic attacks will be feasible.

The term "useful" above is relative to the value of the material being protected by the cryptographic algorithm to the attacker. For example, if the quantum computer attacking a particular key costs US$100 billion to build, costs US$1 billion a year to run, and can extract only one key a year, it is possibly useful to some governments, but probably not useful for attacking the TLS key used to protect a small mail server. On the other hand, if later a similar computer costs US$1 billion to build, costs US$10 million a year to run, and can extract ten keys a year, many more keys become vulnerable.

@@@@@ If the following is wrong, it would be great to have references to replace this with

To date, few people have done systematic research that would give estimates for when useful quantum-based cryptographic attacks might be feasible, and at what cost. Without such research, it is easy to make wild guesses but those are not of much value to people having to decide when to start using post-quantum cryptography.

For example, in [NIST8105], NIST says "researchers working on building a quantum computer have estimated that it is likely that a quantum computer capable of breaking 2000-bit RSA in a matter of hours could be built by 2030 for a budget of about a billion dollars". However, the referenced link is to a YouTube video [Mariantoni] where the researcher, Matteo Mariantoni, says "maybe you should not quote me on that". [NIST8105] gives no other references for predictions on cost and availability of useful cryptographic attacks with quantum computers.

5.1. Proposal: Public Measurements of Various Quantum Technologies

In order to get a rough idea of when useful cryptographic attacks with quantum computers may be feasible, researchers creating such computers can demonstrate them when they can break keys a quarter the size of those in common use. That is, given that 2048-bit RSA, 256-bit elliptic curve, and AES-128 are common today, when a research team has a computer than can break 512-bit RSA, 64-bit elliptic

curve, or AES-128 where only 32 bits are unknown, they should demonstrate it.

Such a demonstration could easily be made fair with trusted representatives from the cryptographic community using verifiable means to pick the keys to break and verifying the time that it takes to break each key. It might be interesting to run the same tests in classical computers at the same time to give perspective.

Note that this proposal would only give an idea of how public progress is being made on quantum computers. Well-funded military agencies (and possibly even criminal enterprises) could be way ahead of the publicly-visible computers. No one should rely on just the public measurements when deciding how safe their keys are against quantum computers.

6. IANA Considerations

None, and thus this section can be removed at final publication.

7. Security Considerations

This entire document is about cryptography, and thus about security.

See Section 1.1 for an important disclaimer about this document and security.

This document is meant to help the reader predict when to transition from using classical cryptographic algorithms to post-quantum algorithms. That decision is ultimately up to the reader, and must be made not only based on predictions of how quantum computing is progressing but also the value of every key that the user handles. For example, a financial institution using TLS to protect its customers' transactions will probably consider its keys more valuable than a small online store, and will thus be likely to begin the transition earlier.

8. Acknowledgements

Some of the ideas here come from Tomofumi Okubo. [[By the time this is finished, this list should be a lot longer.]]

9. References

9.1. Normative References

[Grover96] Grover, L., "A fast quantum mechanical algorithm for database search", 1996, <https://arxiv.org/abs/quant-ph/9605043>.

[Shor97] Shor, P., "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer", 1997, <http://epubs.siam.org/doi/pdf/10.1137/S0097539795293172>.

9.2. Informative References

[Mariantoni] Mariantoni, M., "Building a Superconducting Quantum Computer", 2014, <https://www.youtube.com/watch?v=wWHAs--HA1c>.

[NielsenChuang] Nielsen, M. and I. Chuang, "Quantum Computation and Quantum Information, 10th Anniversary Edition", ISBN 97801-107-00217-3 , 2010.

[NIST8105] Chen, L. and et. al, "Report on Post-Quantum Cryptography", 2016, <http://nvlpubs.nist.gov/nistpubs/ir/2016/NIST.IR.8105.pdf>.

[QuantumSearch] Grover, L., "From Schrodinger's Equation to the Quantum Search Algorithm", 2001, <https://arxiv.org/abs/quant-ph/0109116>.

[RFC3766] Orman, H. and P. Hoffman, "Determining Strengths For Public Keys Used For Exchanging Symmetric Keys", BCP 86, RFC 3766, DOI 10.17487/RFC3766, April 2004, <http://www.rfc-editor.org/info/rfc3766>.

Author's Address

Paul Hoffman
ICANN

Email: paul.hoffman@icann.org