Bell Labs Technical Journal

Volume 23 I November 2018

NOKIA Bell Labs

Quo Vadis Qubit?

Security in the quantum computing era

Dimitrios Schinianakis and Enrique Martín-López
Nokia Bell Labs

Dimitrios Schinianakis (IEEE SM '18) received his diploma in Electrical and Computer Engineering in 2005 and his PhD in 2013, both from the University of Patras, Greece. His current research interests include 5G/IoT security architectures, cryptography, post-quantum and homomorphic encryption, and computer arithmetic. Since 2008, in parallel with his PhD studies, he has been employed in Nokia, first as a Product Line Manager and then as Senior Security Researcher, focusing on IoT and 5G security strategies. He holds several patents with Nokia Bell Labs. He is the recipient of a paper award from IEEE and has delivered tutorials in international conferences on computer arithmetic systems in cryptography. He was an invited speaker in ISCAS 2014 in Melbourne on "Alternative Number Representation Systems". He is an author or co-author of more than 25 papers in international conferences, journals and magazines. He is also a co-author of the books "Secure System Design and Trustable Computing" and "Embedded Systems Design with Special Arithmetic and Number Systems", both by Springer-Verlag. Dr. Dimitrios Schinianakis is a Senior Member of the IEEE and a member of the Technical Chamber of Greece.

Enrique Martín-López received his MS (2010) from the Universidad de Salamanca, Spain, and his PhD from the University of Bristol in 2014, both in physics. Following his doctoral studies in quantum computation and quantum communication using photonics, he joined Nokia to work on quantum technologies. Since then, his research interests have expanded to include distributed consensus protocols, machine learning and artificial intelligence. He has co-authored several scientific publications, which include articles in Science, Nature and Physical Review Letters. He has presented at several international conferences and holds over 20 patents and pending patents. Dr. Martín-López is currently a Research Scientist at Nokia Bell Labs in Cambridge, UK.

Imagine that in a few years from now a full-scale, practical quantum computer hits the headlines. In this apocalyptic scenario, the world of cryptography would be in a state of shock, since almost everything that forms the foundations of current security would collapse. Indeed, the presence of a quantum computer would render state-of-the-art, public-key cryptography useless. All the underlying assumptions about the intractability of mathematical problems that offer confident levels of security today would no longer apply in the presence of a quantum computer.

But are we really doomed? Is cryptography dead? Well, luckily no. This paper examines the technologies that will enable crypto to survive the post-apocalyptic world of quantum computing. There are many things yet to be done to offer an environment as safe as current crypto does, but the tools are there. It is now a game of engineering, pro-active standardization and ingenious mathematics, while following a careful development approach to pave a safe way through the *qubits inferno*.

Introduction

Quantum computing no longer exists solely/exclusively in the wild imagination and dusty drawers of some "unconventional" scientists. Considerable research effort supported by enormous corporate and government funding is currently underway for the development of practical quantum computing systems. The existence of a quantum computer would mark a cornerstone in mankind's technological evolution. It would mean that some computational problems, currently considered intractable for conventional computing systems, would become tractable.

Unfortunately, the intractability of some of these computational problems is the foundation of state-of-the-art cryptography; hence if a practical quantum computer were to be developed, much of today's public-key cryptography infrastructure would need to be replaced by algorithms that can offer the same, if not better, levels of security and resistance against cryptanalysis carried out by quantum computers.

This paper contains a gentle introduction to quantum computing principles and examines the fundamental problems in realizing a quantum

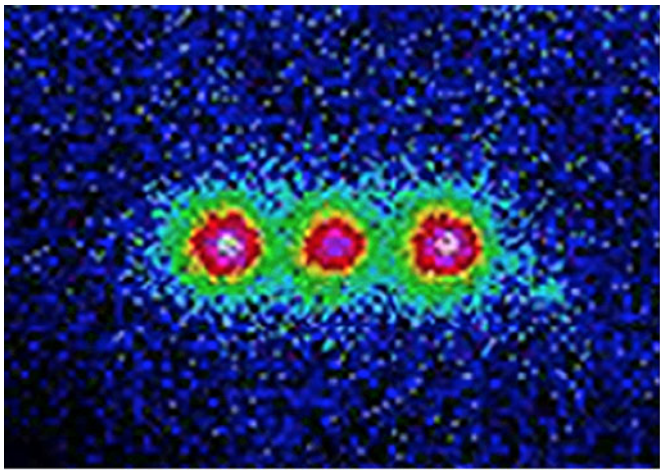

FIGURE 1. Trapped beryllium ions storing three qubits (NIST 2005).

computer. We show how parts of current cryptography would fail in the presence of such a machine. Much to everyone's relief, modern cryptography does offer alternative tools to realize post-quantum cryptosystems; we present these solutions and highlight their basic strengths and pitfalls. Finally, we offer an outlook on future research and some suggestions when considering early deployments of post-quantum cryptography.

Let's not get worried though; even if the foundations of current cryptography would collapse if a practical quantum computer appeared in the news tomorrow, we still have the tools to replace weak cryptographic constructions. However, getting to the safe side requires considerable effort in mathematics, engineering, standardization, hardware design, software and protocol development.

Introduction to quantum computing

It is difficult to offer a simple yet intuitive example of how quantum computers work without risking scientific correctness. Quantum computers are not just better, smaller, or faster versions of conventional computers; they are, rather, based on fundamentally different working principles. Scott Aaronson gives a correct, yet concise explanation: "quantum computing is about choreographing a pattern of interference where the paths leading to each wrong answer interfere destructively and cancel out, while the paths leading to the right answer reinforce each other" (Horgan 2016).

The building blocks of a quantum computer are quantum-bits or "qubits". Think of them as the quantum analogy of conventional bits. In figure 1, three beryllium atoms embodying three qubits are packed together with the use of a "trap" (not shown). We could even say that this is a small-scale quantum computer able to handle a certain set of operations on 3 qubits. The key challenge in building such a machine is about how many qubits we can put together before the laws of classical physics emerge. This would wipe out the quantum dynamics required by quantum algorithms in order to beat their classical counterparts.

The holy grail of quantum computing is to build a so-called *universal quantum computer*, a machine that can be programmed to run any quantum algorithm (just as our current computers can run any classical algorithm). The physical architecture of such a formidable machine will likely depend on the platform (e.g., superconducting circuits, ions) but will likely translate into a number of qubits well beyond our current reach. However, for many, the current challenge is putting together a quantum device that can outperform a classical computer in solving some computational

task, in other words, demonstrating actual "quantum supremacy". But before we delve into these challenges, let's see why quantum computers are so powerful.

The power of qubits

Qubits are the basic unit of information in quantum computing. Similar to classical bits being realized using transistors, qubits are implemented in practice using systems whose physical state adheres to the laws of quantum physics, such as photons or ions. This is where the peculiarities of quantum computers originate. Qubits are fundamentally two-level quantum mechanical systems; therefore their state is defined by a complex vector in a Hilbert space of dimension two. It is due to the complex amplitudes of such vectors that quantum interference effects can happen.

This description will appear unclear to the non-expert reader, hence a rough analogy to binary arithmetic is often used to clarify the logic behind qubit computations. While in conventional computers bits can store either a binary value of '1' or '0', qubits may exist in a combination (superposition) of states '0' and '1' at the same time, which is normal and in line with the laws of quantum physics. During a quantum computing calculation, typically following a quantum algorithm, qubits may exist in any of the exponential number of superpositions of these '1' or '0' states.

However, to output a correct result, a selection needs to be made at the end of the process by measuring the physical state of the qubits. The role of observation in determining the end state of the qubits is famously portrayed by Schrödinger's cat paradox, as described by the renowned physicist in 1935. He suggested a thought experiment, which imagines a cat locked in a steel chamber. Whether the cat is alive or dead depends on the state of a radioactive atom and whether it has decayed (emitted radiation) or not. Schrödinger asserted that the cat would be both alive and dead until the box is opened, and the status of the cat is checked (state is observed). Similarly, the exponential number of superpositions of our binary states can co-exist until measured.

There is currently no quantum algorithm that plays chess. Although this may not be the type of problem best suited to quantum computing, chess is a classic pedagogical paradigm to explain complexity and exponential growth. Let us continue the tradition and use it to illustrate quantum superpositions and quantum algorithms.

Let us imagine that the number of ways in which pieces can be arranged legitimately on a chessboard (also known as the number of diagrams) is about 10^{47} or 2^{156}. Also imagine we could list these 2^{156} diagrams and assign a label to each one of them, which we could achieve by using strings of 156 bits. Let us now prepare a *maximal superposition* of 156 qubits. Such a quantum state would actually correspond to a superposition of all 2^{156} possible diagrams with the same probability amplitude (see figure 2A). This means that if we measured this maximal superposition, we would randomly obtain a single 156-bit string, e.g., "111010...110", which maps to the label of a specific chessboard arrangement (for example the positions in figure 2B or 2C).

Nonetheless, we want to perform useful computations using quantum states not just creating gigantic superpositions that can produce a random diagram selection. So, let us imagine that we prepare a quantum circuit by combining quantum logic gates based on an initial configuration of the board (see figure 2B). Let us also imagine that our 156 qubits, initially in a maximal quantum superposition, are fed into this quantum circuit. This transforms their state into a new superposition, in which not all diagrams have the same probability amplitude. In fact, our mysterious quantum circuit has managed to transform the state of the 156 qubits in a way that only three amplitudes are considerably larger than zero (see figure 2C).

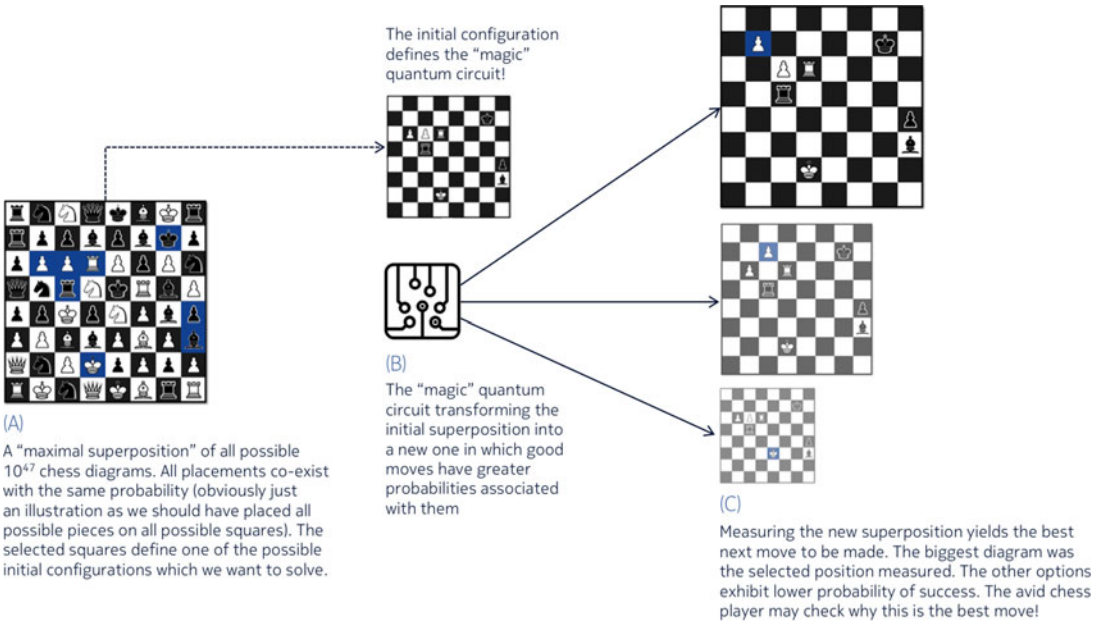

FIGURE 2. A chess analogy to illustrate qubit superposition.

In our chess example, the "magic" quantum circuit acts as good chess players do, instinctively disregarding bad moves and naturally focusing on just a few to decide how to play their next move. Applying wishful thinking once more, imagine that the three amplitudes correspond to three diagrams for the next move with the highest probability of eventually winning the chess game. In such a case, measuring the state of the 156 qubits would randomly obtain a single bit string "101010. . .110", which, with great probability, would correspond to the best move we should finally make (top diagram in figure 2C). The avid chess player may check why the first diagram is the best one!

The property of superposition of states — the fact that a myriad of possible combinations of '1' and '0' states can co-exist together, along with clever logical operations on such states that boost the probability of "solution states" to be measured while producing a destructive interference pattern to the rest— is what allows quantum computers to achieve massive computational gains. For instance, Shor's quantum algorithm for factoring boosts the probability of measuring a number called 'the order', which can be used to reveal the factors of the integer being analysed. **To put it simplistically, while classical computers can only process one possible configuration of the 156-bits of information at a time, quantum computers can account for all possible combinations of those 156-bits simultaneously, due to superposition.** For further reading refer to the excellent textbook of Nielsen and Chuang, Quantum Computation and Quantum Information 2010. Unfortunately, as we shall illustrate, things are not that straightforward when mechanizing a quantum computer.

The decoherence problem

It can be fun to imagine the quantum computation above, but the quantum algorithm behind the "magical" quantum circuit in our chess example has not been (and may never be) discovered. Coming up with quantum algorithms is an extraordinarily hard task, in fact, only a few exist. Interestingly, one of them is Shor's algorithm, the one that can be used to crack public-key cryptography (as we will see shortly). But not only "quantum software" presents a challenge, also "quantum hardware" does. You may think that 156 is a modest number of qubits, but this only refers to the ones required to encode the input of the problem. Many more logical qubits are usually required to create a registry in which the computations can be performed. We have now introduced the term "logical" because, as we will see, a single logical qubit may be formed by a number of "physical" qubits (e.g., ions), which

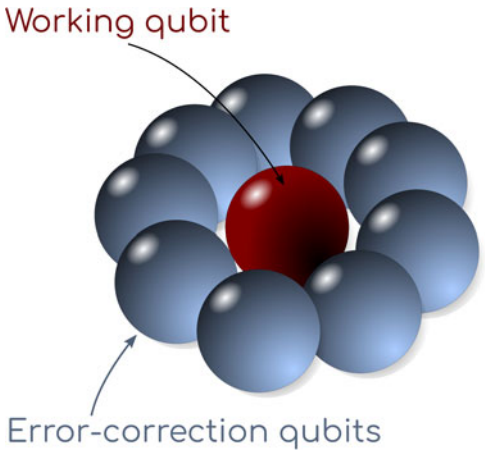

FIGURE 3. A typical structure with one working qubit surrounded by 8 error-correction qubits.

multiplies the required resources. And, as if that wasn't enough, the corresponding quantum gates that make up any required quantum circuitry can be really hard to realize.

There is a physical phenomenon, which makes quantum hardware fundamentally difficult to build, called decoherence. Decoherence is a subtle concept but let us illustrate it again using the Schrödinger's cat experiment. Schrödinger asserted that the cat would be both alive and dead prior to opening the box, but such superposition metaphor could be disrupted by decoherence even if the box was securely closed. The barking of a dog outside the box could make the cat growl. All the effort in keeping the box well shut would be futile as the barking of the dog has acted as a decoherence effect on the superposition state of the cat, collapsing it into "alive". Qubit superpositions are extremely fragile: undesirable interactions with the environment (noise) can act similarly to the measurement stage of a quantum algorithm and destroy the carefully crafted multi-qubit quantum states. Coherence can be understood as the quality of a true quantum superposition of state *amplitudes* rather than a sum of state *probabilities*, which would be the state into which decoherence may drive you. An ensemble of 156 qubits in which some are in state 0 and others are in state 1 with certain probability does not represent a coherent quantum superposition like those required in quantum computing. Therefore much current research is directed at controlling decoherence and eliminating or mitigating noise.

One of the strategies to mitigate the problem is to "couple" a qubit with another qubit, that does not participate in the calculation and can be probed without affecting the main qubit. In principle this means that to construct a functional logical qubit, one would need several physical qubits. A typical structure with one working qubit surrounded by 8 error-correction qubits is illustrated in figure 3. Even with a modest estimation of 800 physical qubits per logical qubit, the overhead is significant (Nielsen, Quantum computing for everyone 2008). An alternative is to cancel-out or to avoid the influence of errors as IBM researchers are investigating (Temme, Bravyi and Gambetta 2017). Figure 4 illustrates qualitatively how errors

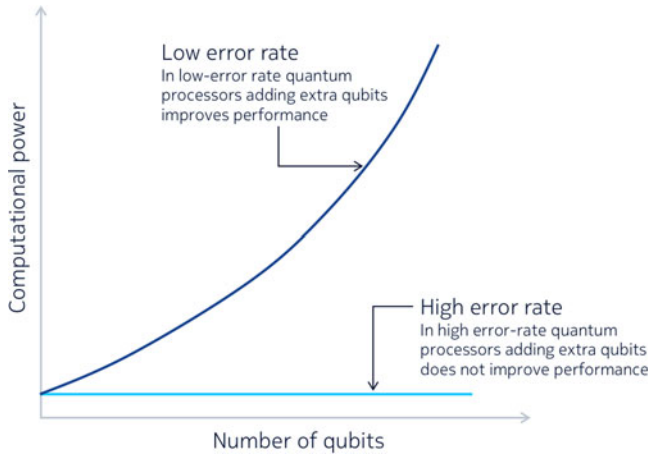

FIGURE 4. Computational power as a function of number of qubits, adopted from (Philip 2018).

affect a quantum processor's performance. Adding extra qubits in a system that exhibits high error rates (high noise) offers virtually no improvements, whereas in low error-rate systems (low noise) the performance increase follows an exponential pattern.

Rather worryingly, there are other researchers who believe that noise control is not just a matter of clever engineering but would eventually lead to a violation of the fundamental laws of computation (Moskvitch 2018). Obviously, if noise cannot be controlled or mitigated, quantum computers will never become a reality. Other scientists suggest that the problem is indeed tangible and that the noise problem boils down to understanding how much of it we can harness.

State-of-the-art

Among others, NOKIA Bell Labs (Nokia Bell Labs 2017), Microsoft, IBM and Intel are all developing their own versions of quantum computing platforms. Google, for example, has rolled out Bristlecone, a 72-qubit processor (Kelly 2018). This new chip does not quite achieve quantum supremacy. It rather serves as a testbed for the research community to investigate error correction techniques, error rates, "as well as applications in quantum simulation, optimization and machine learning". Google is "cautiously optimistic that quantum supremacy can be achieved with Bristlecone".

IBM is planning a 50-qubit platform, while a 20-qubit platform was pipelined for online access by end of 2017 with updates on 2018 (Dignan 2017). Similarly, those platforms don't yet achieve supremacy but serve as valuable tools for experimental, educational and research activities. Along with this, IBM is also developing a software toolkit for quantum application development and execution (Vu 2017). Finally, Intel has unveiled a 49-qubit test processor based on superconducting materials. However, quoting from Intel's website, "it will likely require 1 million or more qubits to achieve commercial relevance" (Intel 2018).

All in all, it is only in the last few years that many of these milestones have been reached in our quest to devise quantum platforms that can really solve problems that conventional computers would need eons to solve. Supremacy is not there yet, but with the available platforms, a wealth of experiments, applications and publications have already been achieved. It is the availability of quantum platforms to the public and research communities that will accelerate the advances in this nascent field.

As a final note, we need to be careful when setting expectations of what a quantum computer could do. Even if quantum supremacy is achieved, the noise control we will have on those machines will still be "imperfect", so we still need to account for those limitations. In fact, state-of-the-art algorithms currently employed by Google to test quantum machines are very specific to benchmarking tasks rather than solving any useful problem such as factoring (Giles and Knight 2018). So far, only a handful of algorithms have been devised that can run on a quantum computer, not to mention that programming in quantum logic is fundamentally different from current software development. Quantum-ready algorithms will take years to mature. In any case, we would still need to come up with new quantum algorithms tackling practical problems that offer a provable advantage over classical computers (quantum machine learning and AI, quantum chemistry simulations for molecular design, etc.). Otherwise, running a quantum computer would be like using a supercar to cross the street.

Cryptography is dead; long live cryptography

Taking a swift look at basic elementary school algebra, it is relatively easy to find the prime factors of 15, namely $15 = 3 \times 5$, but given a large integer, this problem turns out to be intractable for classical computers. This is the problem which one of the most widely deployed cryptographic algorithms, RSA, relies on. The other two most prevalent cryptosystems, Diffie-Hellman key-exchange and Elliptic Curve Cryptography (ECC), employ different hard mathematical problems, the Discrete Logarithm Problem (DLP) and the Elliptic Curve Discrete Logarithm Problem (ECDLP), respectively (Roetteler, et al. 2017).

The accompanying Digital Signature Algorithm (DSA) and the elliptic curve equivalent (ECDSA) for digital signatures are also based on the same mathematical constructions.

Although no formal proof exists that these problems are impossible to solve in a reasonable time frame, no algorithm has been found that can do it since their introduction in the late 1970s. This is, after all, the reason why we gained confidence in these cryptosystems. Unfortunately, this confidence had an expiration date. In 1994, mathematician Peter Shor (Bell Labs researcher at the time), in his seminal work on an algorithm for integer factorization that could be run on a quantum computer, proved that if a large enough quantum computer were available, all modern constructions of public-key cryptography would fail. His work applies directly to the integer factorization problem (RSA), but with minor tweaks, both Diffie-Hellman and ECC constructions would fail as well. As a side note, research has provided estimations on the number of qubits that are required to mechanize Shor's algorithm for large integers. According to the work in (Beauregard 2003) for example, to factor an n-bit integer one needs $2n+3$ logical qubits. This requirement can be reduced by further techniques such as iterative versions of the algorithm requiring just n+1 logical qubits but increased control logic (Parker 2000, Martın-Lopez 2012). Taking (Beauregard 2003) as a reference, for a 2048-bit RSA modulus, one would need 4099 logical qubits, not taking into account the number of extra qubits for error correction. Obviously, this is a number of qubits way beyond the reach of current technology.

Nevertheless, recent announcements by NSA regarding ECC incited a wave of speculation around NSA's cryptanalytic capabilities. On their website, NSA actually proposed that organizations that have not yet adopted ECC should avoid implementing the scheme. Quoting directly from the website:

> For those partners and vendors that have not yet made the transition to Suite B elliptic curve algorithms, we recommend not making a significant expenditure to do so at this point but instead to **prepare for the upcoming quantum-resistant algorithm transition**... Unfortunately, the growth of elliptic curve use has bumped up against the fact of continued progress in the research on quantum computing, which has made it clear that elliptic curve cryptography is not the long-term solution many once hoped it would be. Thus, we have been obligated to update our strategy (NSA 2015).

Though, in the same article, we also read: "· · · it is important to note that we aren't asking vendors to stop implementing the Suite B algorithms and we aren't asking our national security customers to stop using these algorithms".

So, what should the community do?

After the announcement the general impression was that NSA was refraining from using ECC. It is, however, arguable whether NSA can break ECC (Koblitz and Menezes 2015). Nonetheless, the question naturally arises, does the NSA know something more about quantum computers than the rest of the world? Whatever the case may be, this announcement highlights the necessity for research organizations and academia to accelerate their efforts on post-quantum algorithms and protocols (Koblitz and Menezes 2015). Fortunately, things don't look so bad, as the following paragraphs shall illustrate.

Cryptography in the post-quantum era

Before we proceed, an important clarification regarding the term "Quantum Cryptography" needs to be made. The term is often confused with the terms "Quantum-Safe Cryptography" or "Post-Quantum Cryptography". The latter two are essentially the same and are the subject of this paper. Quantum Cryptography, on the other hand, refers to techniques utilizing the laws of physics and quantum mechanical properties to establish secure connections. The most important scheme is Quantum Key Distribution (QKD), which

establishes a common key between two parties. This, in turn, has generated yet another misconception; QKD is not a post-quantum cryptography protocol, but rather a fundamentally secure way to generate a shared secret key that can be subsequently used in symmetric cypher protocols, which, as we will discuss, are believed to be immune to quantum analysis. Quantum Cryptography has great potential but it should not remove the attention from Post-Quantum Cryptography, which offers a more direct solution to quantum cryptanalysis. Moreover, there are important limitations since QKD requires authenticated channels, i.e., the use of traditional cryptography. This somehow contraposes the scope and the problems that quantum cryptography attempts to solve.

Furthermore, while QKD is secure in theory, it doesn't offer real-time information encryption. The quantum channel exhibits noise and attenuation, which must be constrained at acceptable levels in order to maintain a reliable connection. As a result, the current key generation rates and transmission distances are still limited in favor of connection reliability (L. Keuninckx, et al. 2017). This is especially the case in terrestrial applications, where repeaters are required to maintain the quality of the link in acceptable levels. Repeaters add to the complexity of the overall scheme, increase cost and require extra physical security measures to avoid their compromise. An interesting alternative is satellite QKD as has been demonstrated recently by Chinese researchers (Liao, et al. 2018). The link achieved kilohertz-rate key distribution between stations in the cities of Xinglong in China and Graz in Austria, an impressive distance of 7,600 km.

Quantum cryptography is receiving a lot of hype, but the community should not be carried away by its promised capabilities. The associated costs are quite high and not all security considerations for viable deployment have been clarified.

Research has produced some remarkable constructions that are thought to be safe against quantum computers (Bernstein 2017). Let's explore the three main candidates towards a safe quantum world.

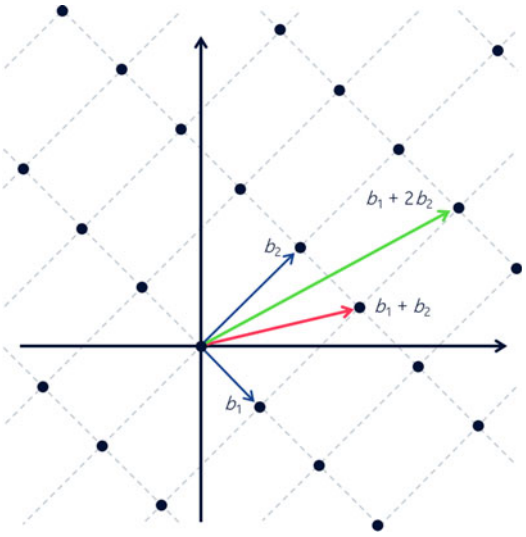

FIGURE 5. An example of a 2-dimensional lattice.

Lattice-based cryptography

Lattices are wonderful mathematical constructions generated by a set of vectors. As seen in figure 5, taking two vectors b_1 and b_2, we can make linear combinations of them to generate new vectors (for example, by simply adding them or multiplying them with integers and then adding them, etc.). By taking **all** (theoretically) possible combinations for the vector sets (b_1, b_2) we can generate a lattice. This grid of points can also include multi-dimensional lattices as well. Just as RSA owes its security to the difficulty of finding the prime factors of an integer, lattice-based crypto (LBC) relies on the difficulty of navigating through the myriad of points in a lattice. Typically, we start with a random point in a multi-dimensional plane and we look for the vector of the lattice closest to our random point. It turns out that for certain families of lattices, this problem is intractable, even for quantum computers (Hoffstein, Pipher and Silverman 1998).

At the beginning, LBC was associated with cryptanalysis failures, but persistent work paid off. The NTRU cryptosystem, for example, has existed for almost 20 years and no solution has been found to break it, even in the presence of

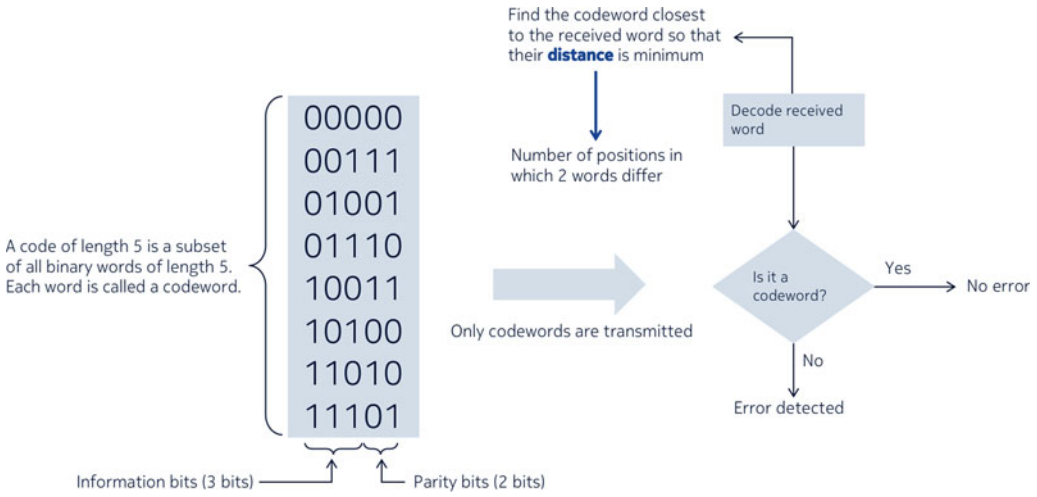

FIGURE 6. The basic principle of binary codes (maximum-likelihood decoding).

quantum platforms (Hoffstein, Pipher and Silverman 1998). The acronym is still a mystery; some believe it stands for "N^{th} degree truncated polynomial ring" and others believe it is "Number Theorists R Us". An interesting fact is that NTRU outperforms RSA and ECC in terms of encryptions/second by orders of magnitude (Hermans, Vercauteren and Preneel 2010). Another major step for fostering research around LBC was the proof that a certain class of lattice-based problems called "Learning With Errors" (LWE) including the more efficient version, Ring-LWE, are secure against quantum computing, as long as the underlying problem of finding the nearest point in a generic lattice is hard for quantum computers (which is the current assumption) (Regev 2005, Peikert 2009, Brakerski et al. 2013, Lyubashevsky et al. 2010, Peikert 2014).

NTRU and Ring-LWE are not as efficient as some of their faster counterparts that have appeared in the literature, but advances in cryptanalysis rendered the latter dangerous for public use. It is said that LBC researchers got a bit greedy while trying to optimize LBC primitives. The only exception is perhaps "New Hope", an optimized Ring-LWE version that avoids certain classes of attacks and promises even better performance (Alkim, et al. 2017). At the time of writing, NTRU, Ring-LWE and New Hope seem to qualify for practical quantum-safe cryptography and early deployments.

Code-based encryption

Code-based encryption is not entirely a new concept. Originally introduced in 1978 by Robert McEliece, it was the first type of encryption that utilized error-correcting codes (McEliece 1978). Its security is based on the problem of decoding a generic linear code. The topic is vast, but figure 6 draws the big picture of error-detection and decoding. In a nutshell, a code of length n is a subset of all binary words of length n. These codewords consist of two parts, namely the information bits and the parity bits. Information bits carry the corresponding information, while the parity bits are used for error-correction. In such a system only codewords are transmitted and during decoding the received words are checked to see if they contain errors. The check is done by finding the codeword that is closest to the received word, so that their distance (bit positions in which they differ) is minimum. If such a codeword exists there are no errors during transmission, otherwise we have detected an error which needs to be corrected. Interestingly, those systems never really got the attention of crypto-practitioners, mainly because the associated private and public keys are quite large (some are MBs in size!) (Augot, et al. 2015).

Code-based encryption exhibits some remarkable properties that make it attractive for the quantum era. The problem of decoding a linear code is an NP-hard problem, so there is no polynomial

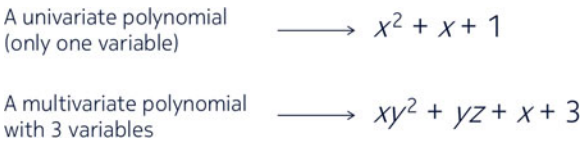

FIGURE 7. Univariate and multivariate polynomials.

time-quantum algorithm known to break the scheme. Of great practical importance is the Niederreiter cryptosystem, a variant of the original McEliece primitive, which allows for short messages, fast encryption and efficient signature schemes. The system was also proven to be resistant against quantum analysis (Niederreiter 1986), (Hang, et al. 2011). Nowadays, there exist efficient software implementations that scale very well, even with RSA standards (Daniel, Chou and Schwabe 2016), which an open research field that could deliver interesting optimizations uses Medium Density Parity Check (MDPC) codes (Misoczki, et al. 2013). These systems tackle some of the problems of long keys and are considered some of the most promising approaches for practical and efficient code-based encryption.

Multivariate cryptography

Multivariate cryptography (MVC) is a relatively new candidate in the field of quantum-safe cryptography. Its security relies on the fact that solving systems of multivariate polynomials (see figure 7) is proven to be NP-hard or NP-complete. An example of multivariate and univariate polynomials is offered in figure 7 for reference. The public keys in these systems are just sets of multivariate polynomials. To encrypt, one simply needs to input the message in these polynomials and perform an evaluation. It's not an oversimplification to say that decryption is just the reverse operation, so one needs to provide the inverse map function for the system of multivariate polynomials. This "trapdoor function" serves as the private key.

There is a great pool of practical signature schemes based on MVC available today, the most prominent being Rainbow (Ding and Schmidt, Rainbow, a new multivariable polynomial signature scheme 2005), UOV (Kipnis, Patarin and Goubin 1999) and pFlash (Ding, Yang, et al. 2007). There also exists the so-called HFEv- family, which produces very short signatures (~120-bits); that are claimed to be much faster than ECDSA, even by today's standards (Petzoldt, et al. 2015). For a full-fledged, public-key encryption scheme, things are still a bit cloudy, because although current schemes allow for fast encryption and decryption, they are also susceptible to decryption failures with a non-negligible probability.

This is not an exhaustive list of quantum-safe options. Many other solutions exist such as primitives based on supersingular isogenies of elliptic curves such as SIKE and SIDH (Costello, Longa and Naehrig Advances in Cryptology, 2017) or hash-based signatures such as XMSS (Buchmann, Dahmen and Hülsing 2011) as alternatives to number theoretic solutions. A high-level snapshot is shown in table 1.

What about symmetric ciphers?

All schemes presented previously — either breakable or unbreakable by a quantum computer—share a common property; they are asymmetric. This means that for two parties to establish a connection there is no need to pre-agree to a common symmetric key. Instead, each party generates its own private and public key-pairs that are subsequently used to establish a common (secret) symmetric key. In its simplest form, Alice could generate a symmetric key, encrypt it with Bob's public key, and send that to him. This secret key can then be used in symmetric ciphers, which are more efficient in encrypting/decrypting big messages than asymmetric ones (fast payload encryption). The prominent AES symmetric algorithm is familiar probably even to noncrypto experts.

Differently from asymmetric ciphers, symmetric ones are not directly threatened by quantum computing (Bernstein 2009). An increase in key

Scheme[1]	Algorithm	Encryption	Signatures	Keys (bytes)	Remarks
Lattice-based	NTRUEncrypt, NTRUSign	✓	✓	1.495 – 2.062	Fast, short keys and signatures, 20 years of analysis
Code-based	McEliece, McBits	✓	✓	~958.482–1.046.739	Large keys, fast, 30-40 years of analysis
Multivariate	Rainbow, Gui, HFEv-	✓ [2]	✓	500.000- 1.000.000	Short signatures, more analysis required
Hash-based [3]	XMSS SPHINCS (Bernstein, et al. 2017)		✓	64 1056	High confidence, large signatures, state management
Elliptic Curve isogenies	SIKE, SIDH	✓	✓	564	Smallest keys and resources but no confidence yet

1. The table contains indicative examples and general features. Details can be found in the references
2. Extra effort to tackle decryption errors is required.
3. An IETF internet-draft for hash-based signatures is already available (McGrew and Curcio 2014)

TABLE 1. Overview of post-quantum cryptographic solutions.

lengths (for example more than 256 bits) will be required as more computing power becomes available in general. This is, in any case, a safe security precaution even by today's computing standards.

Conclusions and outlook

At this point, there are more open questions than answers about the future of quantum computing. If the laws of physics do not pose a fundamental barrier for scalable quantum computing, a powerful enough embodiment of such machine could be expected within the next 5-30 years. After all, NIST has already called for proposals on quantum-safe key exchange and signatures (NIST-CSRC 2017), while IETF and ETSI are also promoting standardization of post-quantum schemes. 3GPP is also actively discussing similar topics in TR 33.841 (Meredith 2018). The PQCrypto conference series is an interesting source for state-of-the-art advances in the field (PQCRYPTO 2018).

In terms of availability, we will not likely have a quantum computer on our desks. We envision that such platforms will be in enterprise, educational/ research or government organizations that can sustain the required resources. Following the current trend of cloud services, processing time on these machines will likely be leased to end users who would like to run their applications on these platforms. This is offered by Rigetti Computing or IBM Q Experience.

A huge concern is the deprecation strategy of current cryptography. Obsolete cryptographic primitives are still an issue, so if a quantum computer were to be released in a few years' time, managing and discontinuing weak crypto would not be a trivial task. Moreover, backdoors and vulnerabilities are still being published for mature schemes. As a result, it is safe to assume that new algorithms will also be found prone to weaknesses. This would call for immediate corrections, so alternatives need to be ready for deployment by that time.

Finally, there is always the concern of "store now - decrypt later". Adversaries with adequate resources, who can intercept and store sensitive data today, might be able to decrypt them as soon as quantum computing becomes available. This is just another argument for early awareness of post-quantum cryptography and speedy deployment of solutions.

On the practical side, our analysis suggests that lattice-based systems and elliptic curve isogenies could constitute the first practical, post-quantum scheme deployments. After all, lattice-based systems are already part of TLS protocols and IKE. Thanks to their small keys and speed, they are suitable for constrained environments and fast communications, but cryptanalysis needs to step up to increase confidence in these systems, especially for EC isogenies. Code-based solutions are also nice candidates since they are already quite mature and exhibit high speeds. Even

though their keys are large, they could still be acceptable for applications that can host them and do not require frequent key-exchanges.

Other interesting schemes are hybrid combinations of existing crypto-algorithms with post-quantum ones. Google, for example, has implemented a hybrid of the New Hope (lattice-based) protocol (Alkim, et al. 2017) for the Chrome browser in combination with ECDSA to test the robustness of New Hope in quantum attacks. In case things go wrong, a fallback to elliptic curve crypto is provided.

When it comes to implementation, we see great prospects for engineering and innovation. Industry needs to understand and encompass all the latest advances in this vibrant field. This will allow enterprises to develop safe interfaces between pre- and post-quantum schemes and tackle hardware/software compatibility issues that will definitely arise. The landscape around constrained IoT devices seems quite challenging, since most of the primitives we examined require large keys, which suggests significant hardware/software resources. Apparently, key players in hardware security need to renew their portfolio of safe security modules (TPM/TEE/HSM) or tailored crypto-processors to support post-quantum crypto.

A quantum-world is definitely not as intimidating as one might expect from a cybersecurity point of view. As this paper has shown, we have solid schemes with which to work, and crypto-experts around the world are pushing for the best possible solutions. It is important to be aware of what quantum computers can and cannot do. They will obviously not be able to make every problem tractable from the outset, and clarifying the vulnerabilities that they may introduce in current cryptography is probably the best starting point in evolving towards a quantum safe era.

References

Alkim, Erdem, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. 2017. "Post-quantum key exchange - a new hope." https://eprint.iacr.org/2015/1092.

Augot, Daniel, Lejla Batina, Daniel J. Bernstein, Joppe Bos, Johannes Buchmann, Wouter Castryck, Orr Dunkelman, et al. 2015. *Initial recommendations of long-term secure post-quantum systems.* PQCRYPTO.

Beauregard, Stephane. 2003. "Circuit for Shor's algorithm using 2n+3 qubits." https://arxiv.org/pdf/quant-ph/0205095v3.pdf.

Bernstein, Daniel J., Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Michael Schneider, Peter Schwabe, and Zooko Wilcox-O'Hearn. 2017. *SPHINCS: practical stateless hash-based signatures*. Accessed July 23, 2018. https://sphincs.cr.yp.to/.

Brakerski, Zvika, Adeline Langlois, Peikert Chris, Regev Oded, Stehlé Damien. 2013. "Classical Hardness of Learning with Errors". *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing. STOC '13*. New York, NY, USA: ACM: 575–584. arXiv:1306.0281. doi: 10.1145/2488608.2488680.

Buchmann, Daniel J. and Tanja Lange. 2017. "Post-quantum cryptography." *Nature*, 549: 188–194.

Buchmann, J., E. Dahmen, and A. Hülsing. 2011. "XMSS — A Practical Forward Secure Signature Scheme Based on Minimal Security Assumption." PQCrypto, Lecture Notes in Computer Science, Springer.

Costello, C., P. Longa, and M. Naehrig. Advances in Cryptology. 2017. "Efficient Algorithms for supersingular isogeny Diffie-Hellman." https://eprint.iacr.org/2016/413.

Bernstein, Daniel J., Tung Chou, and Peter Schwabe. 2016. "McBits: fast constant-time code-based cryptography." https://binary.cr.yp.to/mcbits-20130616.pdf.

Dignan, Larry. 2017. *IBM outlines 50 qubit quantum computing prototype*. ZDNet. November 10. Accessed 2018. https://www.zdnet.com/article/ibm-outlines-50-qubit-quantum-computing-prototype/.

Ding, J., and D. Schmidt. 2005. "Rainbow, a new multivariable polynomial signature scheme." *ACNS, Lecture Notes in Computer Science* 3531: 164–175.

Ding, J., B.-Y. Yang, V. Dubois, C.-M. Cheng, and O. Chen. 2007. "Breaking the symmetry: a way to resist the new differential attack." http://eprint.iacr.org/2007/366.

Giles, Martin, and Will Knight. 2018. *MIT Technology Review: Google thinks it's close to "quantum supremacy." Here's what that really means.* March 9. https://www.technologyreview.com/s/610274/google-thinks-its-close-to-quantum-supremacy-heres-what-that-really-means/.

Hang, Dinh, Moore Cristopher, Russell Alexander, and Rogaway Philip. 2011. "McEliece and Niederreiter cryptosystems that resist quantum Fourier sampling attacks." *Advances in cryptology—CRYPTO, Lecture Notes in Computer Science* 761–779.

Hermans, J., F. Vercauteren, and B. Preneel. 2010. "Speed Records for NTRU." *Lecture Notes in Computer Science* 5985.

Hoffstein, J., J. Pipher, and J.H. Silverman. 1998. "NTRU: A ring-based public key cryptosystem." *Lecture Notes in Computer Science* 1423.

Horgan, John. 2016. Accessed June 24, 2018. https://blogs.scientificamerican.com/cross-check/scott-aaronson-answers-every-ridiculously-big-question-i-throw-at-him/.

Intel. 2018. *2018 CES: Intel Advances Quantum and Neuromorphic Computing Research.* January 08. https://newsroom.intel.com/news/intel-advances-quantum-neuromorphic-computing-research/.

Kelly, Julian. 2018. *A Preview of Bristlecone, Google's New Quantum Processor.* Google Research Blog. March 05. https://research.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html.

Keuninckx, L., M. C. Soriano, I. Fischer, C. R. Mirassoi, R. M. Nguimdo, and G. Van de Sande. 2017. "Encryption key distribution via chaos synchronization." *Scientific Reports* 64 (2): 10–12.

Keuninckx, Lars, Miguel C. Soriano, Ingo Fischer, Claudio R. Mirasso, Romain M. Nguimdo, and Guy Van der Sande. 2017. "Encryption key distribution via chaos synchronization." *Scientific Reports* 64 (2): 10–12.

Kipnis, A., J. Patarin, and L. Goubin. 1999. "Unbalanced oil and vinegar schemes." *EUROCRYPT, Lecture Notes in Computer Science*. 206–222.

Koblitz, N., and A. J. Menezes. 2015. "A riddle wrapped in an enigma." https://eprint.iacr.org/2015/1018.pdf.

Liao, Sheng-Kai, Wen-Qi Cai, Johannes Handsteiner, Bo Liu, Juan Yin, Liang Zhang, Dominik Rauch, et al. 2018. "Satellite-Relayed Intercontinental Quantum Network." *Phys. Rev. Lett.* 120 (3).

Lyubashevsky, Vadim, Peikert Chris, Regev Oded. 2010. "On Ideal Lattices and Learning with Errors over Rings". *Advances in Cryptology—EUROCRYPT*. Springer, Berlin, Heidelberg: 1–23. doi:10.1007/978-3-642-13190-5_1.

Martín-López, E., et al. 2012. "Experimental realization of Shor's algorithm using qubit recycling". *Nat. Photon.* 6: 773–776. https://www.nature.com/articles/nphoton.2012.259

McEliece, Robert J. 1978. *A public-key cryptosystem based on algebraic coding theory.* Jet Propulsion Laboratory. http://ipnpr.jpl.nasa.gov/progress_report2/42-44/44N.PDF.

McGrew, D., and M. Curcio. 2014. "Hash-Based Signatures. Request for Comments." Internet Engineering Task Force.

Meredith, John M. 2018. *Study on supporting 256-bit algorithms for 5G*. Accessed July 23, 2018. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId = 3422.

Misoczki, R., J.-P. Tillich, N. Sendrier, and P. S. L. M. Barreto. 2013. "MDPC-McEliece: New McEliece variants from Moderate Density Parity-Check codes." *IEEE International Symposium on Information Theory (ISIT)*. 2013. 2069–2073.

Moskvitch, Katia. 2018. *Quanta Magazine*. February 7. https://www.quantamagazine.org/gil-kalais-argument-against-quantum-computers-20180207/.

Niederreiter, H. 1986. "Knapsack Type Cryptosystems and Algebraic Coding Theory." *Problems of Control and Information Theory* 15.

Nielsen, Michael. 2008. *Quantum computing for everyone*. August 28. Accessed April 24, 2018. michaelnielsen.org/blog/quantum-computing-for-everyone/.

Nielsen, Michael, and Isaac Chuang. 2010. *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.

NIST. 2005. *Wikipedia*. May 3. Accessed 2018. https://commons.wikimedia.org/wiki/File:Fourier_Transform,_Beryllium_Ions_(5883951063).jpg.

NIST-CSRC. 2017. *Post-Quantum Cryptography*. Accessed 6 25, 2018. csrc.nist.gov/groups/ST/post-quantum-crypto.

Nokia Bell Labs. 2017. *Quantum Computing Using Novel Topological Qubits at Nokia Bell Labs*. https://www.nature.com/nature/outline/quantum-computing/pdf/quantum-computing.pdf?

Parker S. and M. B. Plenio, 2000. "Efficient factorization with a single pure qubit and logN mixed qubits." *Phys. Rev. Lett.* 85: 3049–3052. https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.85.3049

Peikert, Chris. "Public-key Cryptosystems from the Worst-case Shortest Vector Problem: Extended Abstract". *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing. STOC '09*. New York, NY, USA: ACM: 333–342. doi: 10.1145/1536414.1536461.

Peikert, Chris. 2014. "Lattice Cryptography for the Internet." *IACR*. Retrieved 2017-01-11.

Petzoldt, Albrecht, Ming-Shing Chen, Bo-Yin Yang, Chengdong Tao, and Jintai Ding. 2015. "Design Principles for HFEv- Based Multivariate Signature Schemes."

Philip, Ball. 2018. *Quanta Magazine*. January 24. Accessed April 2018. https://www.quantamagazine.org/the-era-of-quantum-computing-is-here-outlook-cloudy-20180124/.

PQCRYPTO. 2018. *ICT-645622*. Accessed 6 25, 2018. https://pqcrypto.eu.org/index.html.

Regev Oded. "On Lattices, Learning with Errors, Random Linear Codes, and Cryptography". *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing. STOC '05*. New York, NY, USA: ACM: 84–93. doi: 10.1145/1060590.1060603.

Roetteler, Martin, Michael Naehrig, Krysta M. Svore, and Kristin Lauter. 2017. "Quantum Resource Estimates for Computing Elliptic Curve Discrete Logarithms." https://eprint.iacr.org/2017/598.pdf.

Temme, Kristan, Sergey Bravyi, and Jay M. Gambetta. 2017. "Error mitigation for short-depth quantum circuits." *Physical Review Letters* 119 (18).

Vu, Christine. 2017. *IBM Announces Advances to IBM Quantum Systems & Ecosystem*. IBM. November 10. http://www-03.ibm.com/press/us/en/pressrelease/53374.wss.