We offer Denmark's strongest and most versatile competences within consultancy, development and application of information technologies

Whitepaper

Post-Quantum Cryptography

Tore Kasper Frederiksen

July, 2018

ALEXANDRA INSTITUTE

ALEXANDRA INSTITUTTET

1 Introduction

L	R	O
0	0	0
0	1	0
1	0	0
1	1	1

Figure 1: Truth table for a classical AND gate.

Our classical notion of nature and computing consists of things happening in a deterministic manner; two plus two is always four. But what if it is mostly four, but once in a while it is three, or maybe five? Regular computation works on bits, which takes the value of either 0 or 1. Two bits are then given as input to a small physical device, a *gate*, which outputs another bit whose value depends on the value of the two input bits. This is always deterministic and can thus be expressed in a *truth table*, such as in Figure 1. However, with the advent of quantum theory, our understanding of the universe in general went from it being deterministic to it being probabilistic; shaped by probability functions. It wasn't until the 1980's that this model of the physical world carried into the world of computation with the field of *quantum informatics*. In this worldview we do not have bits but *qubits*, which take a value of 0 with some probability and a value of 1 with some other probability. Quantum gates then take a list of qubits as input and return a list of qubits as output. This has led to a completely different way of thinking as the probabilities of the different qubits can be entangled with each other through the quantum gates. This leads to a quite different model of computation, and it turns out that there are certain problems that can be solved efficiently on a quantum computer, which currently take a very long time to solve on a classical computer. Even so, it has not been proven whether the quantum computer is inherently more powerful than the classical one, or if we simply have had an easier time coming up with algorithms in the quantum model, which have yet to be discovered counterparts in the classical model.

One might think that the problems we *know* how to solve efficiently on a quantum computer, but don't *know* how to do efficiently on a classical one, are only of theoretical interest. Unfortunately, that is not the case. Many of these problems are directly linked to the problems whose hardness we rely on for security in many cryptographic schemes. In particular, this means that there are schemes, for example RSA or ElGamal, that can be broken quickly on a quantum computer.

Fortunately, quantum computers are currently only able to handle a two-digit amount of qubits and so keys of hundreds or thousands of bits are still safe. However, research moves quickly, so it is likely that it won't be too long before the schemes we rely on today fall victim to the quantum computer.

1.1 Quantum Attacks

The area of cryptography which has suffered the most under the quantum computer is asymmetric cryptography, such as RSA and ElGamal, along with Diffie-Hellman key exchange (both standard and the elliptic curve approach). Still, symmetric key cryptography and hash functions are not free from issues with quantum computers in the world. Two algorithms are to blame for classical cryptography's problems. These are Shor's algorithm [46] and Grover's algorithm [23].

Shor's algorithm allows a quantum computer to solve the discrete algorithm *and* the factorization problem efficiently. Thus, popular schemes such as RSA, DSA, ElGamal, and Diffie-Hellman key exchange become broken as soon as quantum computers are able to handle a few thousand qubits.

Grover's algorithm on the other hand, is very general and works in many settings. Basically, it greatly reduces the time it takes to do a brute-force search. For symmetric cryptography it concretely means that we must double key lengths to keep our current level of security in the post-quantum world. Thus, if we are currently satisfied with the security offered by a κ bit key against a classical computer, then we will get similar security against a quantum computer if we use a 2κ bit key. However, it should be noted that Grover's algorithm cannot be parallelized (unlike a classical brute-force search) so that each iteration of Grover's algorithm *must* finish before the next iteration can start.

It is also worth noting that there is an area of cryptography where we are *guaranteed* that quantum computers won't have any advantage. That is the area of *information theoretic cryptography*, such as the one-time-pad.

1.2 Post-Quantum Cryptography

The academic (and even commercial space[1]) has not taken the threat of quantum computers lightly, and a lot of work has been carried out to develop schemes conjectured to be secure against attacks of quantum computers. Most of these schemes are based on lattices. However, some schemes have also been made based on traditional hash functions, multivariate quadratic equations and isogeny of elliptic curves.

Furthermore, even government agencies are starting to take the quantum threat seriously. In particular we note that standardization work is being carried out by NIST for post quantum algorithms (see https://csrc.nist.gov/projects/post-quantum-cryptography) and we strongly recommend taking this standardization work into account when selecting post-quantum schemes.

1.3 Quantum Cryptography

The special properties of the quantum computer is not only a curse for cryptographers, but also a blessing. The unique quantum-based model has led to the development of some interesting cryptographic schemes with features, such as tamper detections, which are not normally possible to achieve to the same extent in the classical setting. This field is known as *quantum cryptography* and thus distinguishes itself from post-quantum cryptography by requiring quantum mechanics in the schemes themselves, whereas post-quantum cryptography runs on classical computers but tries to thwart attacks made by quantum computers.

[1] See, for example www.ntru.com.

1.4 Outline

In the following sections we go through the different cryptographic primitives in use today in the commercial setting, and discuss how they can be made post-quantum secure or what the best post-quantum secure alternatives are. We do this in two sections; one considering *Non-number Theoretic* cryptography and another considering *Number Theoretic* cryptography. The non-number theoretic section looks into what we can do under the assumption that *one-way* functions exist. That is, we assume there exist functions which are easy to compute, but hard to inverse. This is sufficient to cover symmetric encryptions, Message Authentication Codes (MACs) and hash functions. The number theoretic section adds the assumption that *trapdoor* one-way functions exist (based on number theoretic assumptions). This means that we assume it is possible to construct a one-way function with an auxiliary piece of information, the trapdoor, making it possible to inverse the function. This covers public key cryptography, signature schemes and key exchange.

2 Glossary

Assumption: In the cryptographic sense, an assumption is a mathematical problem, or property, believed to hold, but which has not been formally proven to hold. We wish assumptions to be as weak as possible (meaning that they are easy to believe and have a lot of credence).

CCA: Chosen Ciphertext Attack. A standard and strong definition of security of public key encryption.

Forward Secrecy: The compromise of a long-term secret key does not compromise previous sessions where this key was used.

LWE: The Learning With Errors. A problem which many lattice-based constructions reduce to. The problem is conjectured to be hard to solve using a quantum computer.

NP-hard: A complexity class defining the set of problems which we believe requires a long time to solve, but where a solution can be verified efficiently.

Side-channel attack: An attack of a cryptographic scheme based on observing the execution of an implementation, rather than by trying to break it based on weaknesses in the scheme itself.

Timing attack: An specific type of side-channel attack leveraging that the execution of a certain operation depends on the value of the (secret) input to this operation.

3 Non-number Theoretic Cryptography

3.1 Symmetric Key Cryptography

Both block ciphers and stream ciphers are affected by Grover's algorithm. Thus we recommend doubling the of key-size over what is acceptable to thwart classical attacks in order to gain conjectured security against quantum computers.

3.2 Hash Functions

If we consider the standard requirements of hash functions, i.e. hard to find a preimage, hard to find a second preimage and hard to find a collision, then we have the same case as for symmetric encryption schemes and must double the key length. There are, however, a few caveats. It has been shown that quantum attacks for breaking collision resistance can be improved slightly over Grover's algorithm. More concretely it is required to increase the digest size by an order of 2.5 [10], rather than 2 (as for the symmetric schemes). An older result shows that the increase must rather be a factor 3 [9]. However, this is illusionary as it assumes an unrealistically large quantum memory of the adversary. Still, because this is an active area of research, where improvements are continuously found, it might be a good idea to err on the side of caution and increase digest length by a factor 3 as a minimum.

3.3 MACs

A MAC is the secret key equivalent of a digital signature, meaning that the same key is used for both signing and verification. Such schemes can for example be based on block ciphers (CBC-MAC) or hash functions (HMAC). For such constructions we note that, in the worst case, one needs to increase key lengths by the same factor as their underlying primitives. We say at worst, since finding a collision on the underlying hash function does not necessarily mean that HMAC can be broken.

3.4 Conclusion

The recommendations in this section are first of all based on the currently best known quantum attacks and do not consider the stronger setting where the adversary holds some auxiliary quantum information it can use in its attack. In such attacks, where the adversary can query an oracle in some maliciously chosen quantum state, it turns out that many modes of operation for MAC schemes can be broken easily [27].

Furthermore, it turns out that there are problems, relevant to symmetric cryptography, that can be solved efficiently using a quantum computer [32]. However, it is not clear how to use these to perform general and efficient attacks on symmetric encryptions or hash functions. Still, in time they may be developed into algorithms giving significantly better attacks than simply brute-force.

If one is willing to assume that an adversary can access *classically* encrypted values or digests then everything suggests that increasing the security parameter with a constant factor will be sufficient to thwart attacks. Currently this factor is 2, except for finding hash collisions where it is 2.5.

4 Number Theoretic Cryptography

Shor's algorithm [46] showed that using a quantum computer, one can easily factor a large integer into its prime components. He likewise showed that computing the discrete logarithm could be done efficiently as well. This result makes basically all public key encryption schemes, digital signature schemes and key exchange algorithms in use today vulnerable to quantum attacks, including but not limited to RSA, ElGamal, (EC)DSA, and Diffie-Hellman.

We note that many of the post-quantum schemes we suggest below will have "large" key sizes, by which we mean that they will be greater than the key size of the schemes in use today. As a concrete reference we note that RSA and ElGamal only use a few kilobytes for both public and private keys, along with ciphertexts. For schemes based on elliptic curves, it might be less than a single kilobyte.

Post-quantum Assumptions. With a few exceptions, every system we suggest in the following will be based on a mathematical structure known as lattice. It is by far the most pervasively used family of structures that post-quantum cryptography can be based on. However, there is not one clear lattice assumption that will fit everything, so here we go through a bit of background in regard to security and the assumptions lattice cryptography relies on.

At the very high level, a lattice is a mathematical group which is represented by a multi-dimensional grid. Thus an element in this group is a point in the lattice (grid). A lattice can thus be expressed as a a set of vectors of certain dimensions representing the base of this gird. More specifically, one might think of it as an integer matrix. Most schemes based on lattices rely on the hardness of (some variant of) the *Learning With Errors* (LWE) problem, which becomes increasingly hard when the lattice requires many basis vectors in high dimensions. This is a very interesting problem since it has been shown [45] to have a *worst case to average case reduction*. This means that any randomly picked instance of this problem is as hard to solve as the hardest instances of the problem. This is indeed a very positive feature to have in a cryptographic problem as it means that there are no bad random instances. Unfortunately, schemes based on LWE require large keys and are a bit slow. Work has therefore been carried out using other lattice assumptions. Some of these relate to LWE in that the overall problem is the same but with some different structures. This is for example the case for the Ring-LWE problem, which is based on a type of lattice known as ideal. These contain more structure than "standard" lattices and thus *may* open up for more attacks.

Lattices and assumptions “in between” the Ring-LWE and general LWE are based on Module-LWE. This is a generalization of Ring-LWE and places itself between Ring-LWE and general LWE and is thus more desirable than Ring-LWE. Finally, we have non-standard lattice assumptions such as the NTRU assumption. This assumption is tightly bound to specific schemes and does not have a reduction to a standard problem.

For convenience, we relate the assumptions by inequality on desirability: LWE $\geq$ Modulo-LWE $\geq$ Ring-LWE $\geq$ NTRU. We note that equality in this hierarchy means that one thing might not *necessarily* be better than the other. Still, the hierarchy is based on a theoretical view of what we currently know. Thus LWE is more desirable than Modulo-LWE, but they are both more desirable than Ring-LWE and NTRU. It should come as no surprise that the more desirable the assumption, the slower the scheme, and the larger key sizes are in play. Thus picking or designing a scheme often becomes a trade-off between *expected* security and efficiency.

Furthermore, one thing we must be aware of when using lattices is the risk of *side-channel attacks*. These are attacks on the implementation of a scheme rather than the (theoretic) scheme itself. For lattice a specific kind of side-channel attacks known as *timing attacks* are of particular interest. These are attacks where the adversary tries to learn some secret values by observing how long *non-constant* operations of the program take. By non-constant time we mean that the time it takes to execute an operation depends on the value of its input.

This is particularly relevant in the case of schemes based on the LWE problem as they tend to require sampling of a normal distribution over the integers, and most algorithms completing this task don’t run in constant time. Still, simply avoiding this sampling or using constant time algorithms for this task might not be enough. It is crucial that the entire implemented scheme executes in constant time to avoid timing attacks.

Finally, it should be noted that a recent paper claimed to have broken the quantum security of certain types of lattice assumptions [20]. The paper was, however, quickly retracted, as it contained a bug that did not seem possible to fix. Still, one of the authors was Peter Shor, who was behind the efficient quantum algorithms for solving factoring and discrete logarithm, so the fact that he is working, and making progress, in quantum algorithms for lattice problems should motivate one to use caution in regard to these systems.

4.1 Encryption

Several quantum secure alternatives exist which can replace public key encryption schemes such as RSA or ElGamal. The interesting alternatives can be classified into two different families: lattices and error correction codes. Which family, and which specific scheme, will be the best replacement depends highly on the setting of usage, as they all have their pros and cons.

4.1.1 McEliece

The McEliece scheme [36] is old, from well-before Shor's algorithm came along. However, it has started to receive new attention because of its conjectured quantum security. The overall idea of the McEliece scheme (and its follow-up works) is that a linear error correction code is used to encode the plaintext into a code word. Random errors are then added to this code word in such a way that only the party who knows a special trapdoor can remove these errors and decode the encoding. The scheme can be made CCA secure using simple conversions [31]. In fact, this must be done to avoid some attacks on the system. The schemes turn out to be particularly well-suited for computationally constrained devices, since the computations required by the scheme can be carried out using standard bit operations on small words. Unfortunately, the initial version of the scheme has rather large public keys. However, through the years the scheme has been optimized. Of these optimizations, the Niederreiter scheme [40] is of particular interest. This version manages to reduce the public key size to around 200 bytes (for reasonable security requirements). The Niederreiter scheme further has the advantage that it can be adapted into a signature scheme as well [14].

Other schemes based on McEliece exist, for example MDPC-McEliece [38]. These schemes have different security assumptions which have not been studied thoroughly enough to be deemed recommendable by experts in the field [6].

One should however be aware that because of the encoding/decoding procedures involved, it can easily become relatively inefficient to implement McEliece or Niederreiter without being vulnerable to timing attacks.

Pros:

- The scheme has resisted 40 years of scrutiny.
- Computationally efficient (implemented using binary matrix multiplication), in particular suitable for computationally constrained devices.
- Can be made CCA secure.
- Can be adapted to a signature scheme.

Cons:

- Large public keys (several hundred kilobytes).
- The underlying security assumption that the scheme is reduced to is not too well-understood.
- Hard to implement without timing attack vulnerabilities.

If using a code-based encryption scheme, the Niederreiter system [40] seems the most promising.

4.1.2 NTRU

An NTRU scheme first saw the light of day in 1996 and thus represents one of the oldest families of lattice-based crypto schemes. NTRU consists of two schemes; an encryption scheme, NTRUEncrypt, and a signature scheme, NTRUESign. The schemes started out patented and commercialized by the aptly named NTRU Cryptosystems, Inc. However, in 2017 the encryption scheme was made open source and is now in the public domain. NTRU is efficient and does not have large public keys, as is the problem of many other lattice- or code-based cryptosystems. However, the basic NTRU schemes lack a reduction to standard and well-studied problems. Still, a slight variant of NTRU was introduced by Stehlé and Steinfeld [47], which offers a security reduction to a specific variant of Ring-LWE.

Recently, a new variant of NTRU was introduced called NTRU Prime [5]. This scheme removes certain structures present in the NTRU system. Though these structures have not compromised the NTRU system, they have had an influence in compromising other lattice-based systems. The idea is that removing these will increase the security of the scheme. Even so, there is no reduction to a standard lattice problem. Still, the argument can be made that this is not an issue, as the NTRU scheme has been studied thoroughly for many years, close to the extent of the underlying problems other schemes reduce to.

Unfortunately, significant progress has recently been made in attacking the NTRU schemes, for certain choices of parameters [2, 30]. Even against NTRU Prime in certain cases [30]. These attacks cast doubt on the general security of NTRU and requires users to take extreme care when deciding on parameters.

Pros:

- Computationally *and* communicationally efficient.
- Constant time execution.

Cons:

- Plain NTRU has a lot of structures that *could* make it vulnerable.
- Efficient attacks exist on certain choices of parameters.
- No reduction to a “standard” lattice problem.

Of the NTRU variants, the NTRU Prime scheme [5] seems to be the least vulnerable.

4.1.3 LWE

Several schemes based on LWE exist, starting with the foundational work of Regev [45]. His scheme had a few downsides though, such as not being very efficient, having large keys and not considering CCA security. Still, based on the ideas in Regev’s work, several pieces of follow-up work have been completed, culminating in Lizard [11], which managed to keep security reducible to the "standard" LWE assumption. Furthermore, Lizard has fast encryption and decryption, along with fast key generation, all while being CCA secure. Furthermore, this scheme manages to avoid doing sampling from the normal distribution as part of the encryption process, thus decreasing the likelihood of timing attacks.

Pros:

- Can be reduced to a general version of LWE.
- Does not have any potentially "dangerous" structures.
- CCA secure by construction.
- Does not require sampling from the normal distribution during encryption.

Cons:

- Large public keys and large ciphertexts.
- Lizard has not received too much scrutiny yet.

Of the plain LWE schemes, Lizard [11] appears to be the most desirable.

4.1.4 Kyber

A recent lattice-based cryptosystem, submitted for NIST standardization, is Kyber [8]. This scheme is based on module-lattice assumptions, specifically the module-LWE assumption, and has been designed to withstand timing attacks. The keys are a bit on the large side (order of kilobytes), but both encryption, decryption and key generation are very efficient. The scheme is CCA secure and can also be used for key exchange. It furthermore has the advantage of having a signature variant as well [18], making it possible to use the same assumption for all one's asymmetric needs.

Pros:

- Efficient encryption, decryption, and key generation.
- CCA secure by construction.
- Designed to withstand timing attacks.
- Limited "dangerous" lattice structures, and thus attack surface.

Cons:

- Slightly large keys (on the order of kilobytes).
- Relatively new scheme, which have not yet received too much scrutiny.
- Not reducible to the *most* general lattice problems.

4.1.5 Conclusion

Of the schemes we have presented, Lizard [11] is probably the most desirable overall, as it currently *seems* to give the best security guarantees, without asking for too much of a compromise on efficiency and sizes. However, if public keys of hundreds of kilobytes are not viable, or a less conservative view on security is acceptable, then Kyber [8] is an excellent alternative. On the other hand, if speed is of the utmost importance, then NTRU Prime [5] is an option as well. We do note that good encryption schemes based on ring-LWE also exist [34], but in general Kyber seems to offer more for the same "price".

Regarding schemes based on other assumptions; the Niederreiter [40] is the only desirable alternative if one does not wish to use lattices, or is working on a computationally constrained device.

4.2 Signatures

Several different approaches to achieving quantum secure signature schemes exist. Like for encryption, one of these is lattices. Another is multivariate quadratic equations that use special structures between several quadratic polynomials. Finally, one can make schemes based on a regular hash function, using a tree of hash digests.

As is the case for encryption, each of these have their strengths and weaknesses and so there is no one scheme which is best in all situations. Thus, one must make compromises when choosing, keeping the requirements of one's context in mind.

4.2.1 Multivariate-quadratic-equations

Multivariate quadratic equation-based signature schemes are based on the multivariate quadratic polynomial problem, which has been proven to be NP-hard. Basically, the problem is finding a vector of n values s.t. when they are plugged into each of m quadratic polynomials of n variables, they all evaluate to 0. This is believed to be hard in the average case as well as the worst, but there is currently no worst case to average case reduction. Unfortunately, the security of schemes based on this problem does not reduce to this problem itself, but rather a related problem. Thus, the security reductions are not as desirable as the kind used for lattices.[2]

Many variants of these schemes exist, but so do a lot of attacks. The first scheme was the C* scheme, which was introduced in 1988 by Matsumoto and Imai [35]. However, the C* scheme was broken already back in 1995 [41] and the initial version of another scheme, called *Oil and Vinegar*, was broken in 1999 [29].[3] Still, another family of schemes known as HFE appeared already in 1996 [42], based on the work of Matsumoto and Imai. Of these schemes, one particular kind known as HFEV- [13] remains

[2] We note that the lattice problems and the multivariate quadratic equation problems are different and thus it might be the case that lattice constructions get completely broken, yet that multivariate quadratic equation-based schemes remain secure.

[3] A repaired version was introduced later [28].

unbroken. Furthermore, it does not seem vulnerable to any of the cryptanalytic attacks introduced against multivariate quadratic equation systems, at least assuming a good choice of parameters, such as those suggested by Petzoldt *et al.* [44].

Pros:

- Efficient signing and verification.
- Can be implemented using simple operations.
- Small signatures.

Cons:

- Large public keys (around 100 kilobytes).
- Many of the older schemes have been broken.
- No worst-case to average-case reductions.
- Assumption reductions not as general as for lattice.

If one intends to use a scheme based on multivariate quadratic equations, then the HFEv- scheme with parameters as suggested by Petzoldt *et al.* [44] seems to currently be the most sensible choice.

4.2.2 Hash-based

A hash-based signature scheme was first introduced by Lamport in 1975. His scheme is what is known as a *one-time-signature* scheme, meaning that a public/private key pair can only be used for a *single* signature. This is obviously quite inefficient, since we generally want to be able to sign more than a single message per key. However, optimizations exist, for example using a tree-structure known as a Merkle tree [37] which allows one to sign N messages using a public key of a similar size as in the Lamport scheme. However, each signature contains a fresh Lamport signature scheme in itself, meaning that the private key has size around N times that of the standard Lamport scheme. The essence of the Merkle scheme is thus that it allows one to set up a public key, which can be used for many one-time signatures. Several pieces of follow-up work have been done, culminating in the XMSS-MT scheme [25]. This scheme manages to limit the amount of storage needed, both for the public key and private key, along with improving the time it takes to sign and verify to within a few milliseconds. Unfortunately, generating the keys for the system still takes a significant amount of time (on the order of seconds or minutes for schemes allowing up to a million signatures).

It should be noted that a version of XMSS is undergoing standardization at the Internet Engineering Task Force (IETF). Furthermore, the scheme has the nice property (unlike most current widely used signature schemes) that it is *forward secret*. This means that if the private key/state gets compromised, previous signatures remains valid. This comes in effect since the signatures are one-time from an ordered list, thus it is possible to publish the point on the list where the compromise happened and simply reject all signatures after that point.

Pros:

- Fast signing and verification.
- Forward secrecy.
- Based on standard and widely accepted assumptions on hash functions.

Cons:

- Very long key generation time (up to minutes).
- Somewhat large public keys (order of kilobytes).
- Can only be used for a limited, predetermined amount of signatures.

Currently the XMSS-MT scheme [25] seems to provide the best efficiency without compromising on security.

4.2.3 Lattices

Like for encryption, many different lattice-based signature schemes exist, relying on assumptions such as LWE, Ring-LWE, Module-LWE, and NTRU. Even more than that, there are different families of "styles" such as the *hash-and-sign* family or the *Fiat-Shamir* family.

One of the oldest lattice-based signature schemes is based on NTRU, like the case for encryption. This scheme is called NTRUSign [24]. Unfortunately, the NTRU signature scheme was quickly broken towards an adversary seeing as little as 400 signatures [39]. Even its fix was broken again, although now assuming the adversary gets to see 8000 signatures [19]. Furthermore, usage of the NTRU sign algorithm requires a license. So even though repairs have been suggested to the NTRU signature scheme, we cannot recommend using it for these reasons.

However, researchers have carried out more work based on the underlying NTRU assumption (which is considered safe for certain choices of parameters) resulting in protocols different from NTRUSign. Of particular interest are the BLISS schemes [17, 16], which yield small and efficient signatures. Even though they reduce to an assumption related to NTRU, the actual assumption in their proofs is weaker (i.e. more desirable) than the standard NTRU assumption. However, security of BLISS is based on sampling from the normal distribution as part of the signing. Because the techniques used are not deterministic, it has made BLISS, and its follow-up works, vulnerable to side-channel attacks, both cache- [43] and timing-based [21]. For these reasons, we cannot currently recommend usage of BLISS and its related systems without extreme caution being taken during implementation.

A reoccurring design idea for lattice-based signatures is the *hash-and-sign* idea: To sign a message, one hashes it to a value, interpreted as some place in a lattice, and the signature is then a nearby point. The idea was introduced in the GPV scheme [22]; a scheme which enjoys worst case reductions. Unfortunately, the scheme was not too efficient, neither in key sizes, nor signing and verification time. Thus, much research based on this paradigm has been carried out since then, culminating in Falcon (see https://falcon-sign.info). This scheme uses the type of lattices from NTRU (to get efficiency and succinctness), though. Thus the assumption is moved to an NTRU-based assumption. However, these schemes inherently need sampling from a normal distribution and so, like BLISS, potentially makes them vulnerable to timing attacks.

Fiat-Shamir One of the families of efficient lattice-based signatures that still has significant security credence is the protocols based on Fiat-Shamir. The overall idea is that randomness is sampled as part of the signing process, the message and the randomness is hashed, and then that is used to compute a signature. This has the caveat that the signature might be rejected if it does not fit a certain distribution. Thus care must be taken to avoid vulnerabilities to side-channel attacks.

Two good choices of schemes exist in this family; TESLA [4] and Delithium [18]. TESLA is based either on LWE or Ring-LWE [1]. Delithium on the other hand is based on module-LWE [33]. However, Delithium is deliberately designed to not be vulnerable to timing attacks and has been submitted to NIST for standardization. For this reason we consider Delithium the most desirable of these two choices. Furthermore, neither of these approaches use sampling from the integer normal distribution as part of the signing process. This may make them less likely to suffer timing attacks than those schemes who do, such as the hash-and-sign approaches.

Pros:

- Fast signing, verification and key generation.

Cons:

- Somewhat large public keys (order of a few kilobytes).
- Has not received much scrutiny yet.

For the more security conservative user, the LWE version of TESLA [4] (with much care taken to ensure a constant time implementation) seems to be the best choice, and for everyone else Delithium [18] seems to be the better choice.

4.2.4 Conclusion

Overall, the XMSS-MT scheme [25] appears to be the best choice, assuming one does not require the private key to be used for more than a million signatures, and that a key generation time of up to a minute is acceptable. If that is not the case, then using the Fiat-Shamir family of lattice signatures seems to be the better choice. The more security conservative might use the plain TESLA [4], at the price of efficiency, whereas the less security conservative might use Delithium [18].

4.3 Key Exchange

As most key exchange protocols are either based on public key encryption or Diffie-Hellman, these will be rendered insecure against a quantum computer. In the same case as for encryption and signatures, several possible approaches for post quantum key exchange can be found in the literature. Again, one of the main approaches is to use lattices, another is based on a mathematical structure between elliptic curves, called isogeny.

4.3.1 Isogeny

The isogeny-based systems rely on a rather “new” assumption on the hardness of isogeny of supersingular elliptic curves. Basically, it is based on the hardness of computing certain maps between certain types of elliptic curves. The first version of a key exchange scheme based on this type of problem was broken because of an efficient quantum algorithm solving the underlying isogeny problem on a certain type of curves (non-singular) [12]. However, in 2011 a new scheme based on other types of elliptic curves (supersingular) were suggested [26]. These curves are resilient to the fast quantum algorithm, and so still believed to be hard to solve on a quantum computer. Unfortunately, the underlying problem is not known to be NP-hard, nor having worst-case to average-case reductions. Furthermore, a simpler version of the underlying problem is in fact solvable in polynomial time. This, along with a variant being broken and the young age of the scheme, does not give much credence to its security, and so we will not recommend any schemes based on supersingular isogeny.

4.3.2 LWE

Quantum secure (unauthenticated) key exchange was introduced only recently [15], based on the standard LWE assumption. This leads to follow-up works, but only one of these is a viable contender when considering standard lattice assumptions, namely Frodo [7]. This scheme proceeds like the traditional Diffie-Hellman protocol but uses a lattice construction. In particular, this means that Frodo offers forward secrecy and has easy integration into TLS. The authors offer a constant time implementation, yielding overall execution time of around a couple of milliseconds. Communication is on the large side, though (compared to Diffie-Hellman key exchange) in the order of tens of kilobytes. A downside is that it requires a quantum secure signature scheme to achieve quantum secure authentication, which is required in order to get full quantum security. This is unfortunately inherent to unauthenticated key exchange. Still, we note that even if a classical signature scheme, such as DSA, is used, we still get significant security gains. Specifically, if the key exchange scheme is deployed before the existence of quantum computers, then sessions carried out *before* a quantum computer becomes available will remain secure, even when the adversary can use a quantum computer. However, after that point, future sessions might not be secure as the adversary would now be able to use the quantum computer to spoof signatures and thus break authentication.

Pros:

- Fast computation.
- Forward secrecy.
- Constant time (unlikely to be vulnerable to timing attacks).
- Based on a conservative LWE assumption.

Cons:

- Somewhat large communication complexity (around 10 kilobytes)
- Has not received too much scrutiny yet.
- Requires a separate signature scheme for authentication.

4.3.3 Ring-LWE

Several works have been carried out in this area, based on the solutions above for standard LWE, culminating in "A New Hope" [3]. This scheme is similar to Frodo [7], and so it enjoys many of the same features. In particular the scheme proceeds like the traditional Diffie-Hellman protocol, but using a lattice construction over rings. Thus, it offers forward secrecy and has easy integration into TLS. Furthermore, the authors ensure that all their design/implementation choices are made to achieve constant time, and thus protection against timing attacks. In addition, the public parameters can be generated such that it can be verified that they do not contain backdoors. Finally, the protocol has speeds and communication comparable to RSA-based key exchange. Unfortunately, this is also unauthenticated, and so a signature scheme must be added on top.

Pros:

- Fast computation.
- Forward secrecy.
- Constant time (unlikely to be vulnerable to timing attacks).
- Explicit backdoor prevention.

Cons:

- Somewhat large communication complexity (around 10 kilobytes)
- Based on the less conservative Ring-LWE assumption.
- Has not received too much scrutiny yet.
- Requires a separate signature scheme for authentication.

4.3.4 Conclusion

For the more security conservative user, Frodo [7] appears to be the best choice, and for everyone else A New Hope [3].

5 Conclusion

5.1 Advanced Cryptography

There are still several cryptographic primitives we have not covered in this survey. This includes things like zero-knowledge, commitments and secure multi-party computation (MPC). However, active research is still happening in these areas to ensure that they can also be used in a post-quantum world. Fortunately, for some of these we get quantum security for free, as MPC for example can be based purely on information theoretic primitives, which are not vulnerable to quantum computers. Others, such as commitments, can be based on symmetric primitives where we only need to pay a small price to ensure quantum security. However, many protocols still rely on the hardness of the Diffie-Hellman assumption, or RSA. Thus, care must be taken and research must be done if we want to use these systems in a quantum world.

5.2 Advice

As we have seen, most schemes that cannot be fixed by extending key sizes a little bit can be based on lattice assumptions. Since the study of the (quantum) hardness of these problems only became *really* interesting once they were used in cryptographic schemes, it is sensible to act with a certain wariness on these assumptions. In particular when using schemes based on assumptions that have not been of independent mathematical interest previously, and assumptions that add more structure to objects in order to optimize performance. Like in many paths of life, keeping things simple is often the best option, this is also true for the mathematical structures used in cryptography. Using old, tried and true concepts often also yield the safest results (assuming of course the old concepts are not broken).

Furthermore, in regard to cryptography, it does *not* in fact make sense to put all your eggs in different baskets as the weakest link will almost always break the chain. What is meant by this is that one should ideally pick a suite of schemes based on the *same* underlying assumption, rather than picking schemes based on distinct assumptions. An example could be using the Kyber and Delithium family for encryption, signatures and key exchange.

A very concrete advice to keep in mind when looking at non-standardized, non-production-grade implemented crypto schemes is to be extremely cautious of side-channel attacks. In the case of lattice crypto, especially timing attacks. It is crucial to use an implementation where care has been taken to ensure constant time execution. Ideally, this should go as far as protecting against leakage resulting from different accesses in the memory hierarchy.

References

[1] Sedat Akleylek, Nina Bindel, Johannes A. Buchmann, Juliane Krämer, and Giorgia Azzurra Marson. An efficient lattice-based signature scheme with provably secure instantiation. In David Pointcheval, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *Progress in Cryptology - AFRICACRYPT 2016 - 8th International Conference on Cryptology in Africa, Fes, Morocco, April 13-15, 2016, Proceedings*, volume 9646 of *Lecture Notes in Computer Science*, pages 44–60. Springer, 2016.

[2] Martin R. Albrecht, Shi Bai, and Léo Ducas. A subfield lattice attack on overstretched NTRU assumptions - cryptanalysis of some FHE and graded encoding schemes. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I*, volume 9814 of *Lecture Notes in Computer Science*, pages 153–178. Springer, 2016.

[3] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016.*, pages 327–343. USENIX Association, 2016.

[4] Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In Josh Benaloh, editor, *Topics in Cryptology - CT-RSA 2014 - The Cryptographer's Track at the RSA Conference 2014, San Francisco, CA, USA, February 25-28, 2014. Proceedings*, volume 8366 of *Lecture Notes in Computer Science*, pages 28–47. Springer, 2014.

[5] Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Christine van Vredendaal. NTRU prime: Reducing attack surface at low cost. In Carlisle Adams and Jan Camenisch, editors, *Selected Areas in Cryptography - SAC 2017 - 24th International Conference, Ottawa, ON, Canada, August 16-18, 2017, Revised Selected Papers*, volume 10719 of *Lecture Notes in Computer Science*, pages 235–260. Springer, 2017.

[6] Daniel J. Bernstein and Tanja Lange. Post-quantum cryptography - dealing with the fallout of physics success. *IACR Cryptology ePrint Archive*, 2017:314, 2017.

[7] Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! practical, quantum-secure key exchange from LWE. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 1006–1018. ACM, 2016.

[8] Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, and Damien Stehlé. CRYSTALS - kyber: a cca-secure module-lattice-based KEM. *IACR Cryptology ePrint Archive*, 2017:634, 2017.

[9] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. In Claudio L. Lucchesi and Arnaldo V. Moura, editors, *LATIN '98: Theoretical Informatics, Third Latin American Symposium, Campinas, Brazil, April, 20-24, 1998, Proceedings*, volume 1380 of *Lecture Notes in Computer Science*, pages 163–169. Springer, 1998.

[10] André Chailloux, María Naya-Plasencia, and André Schrottenloher. An efficient quantum collision search algorithm and implications on symmetric cryptography. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*, volume 10625 of *Lecture Notes in Computer Science*, pages 211–240. Springer, 2017.

[11] Jung Hee Cheon, Duhyeong Kim, Joohee Lee, and Yong Soo Song. Lizard: Cut off the tail! // practical post-quantum public-key encryption from LWE and LWR. *IACR Cryptology ePrint Archive*, 2016:1126, 2016.

[12] Andrew M. Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *J. Mathematical Cryptology*, 8(1):1–29, 2014.

[13] Nicolas Courtois, Magnus Daum, and Patrick Felke. On the security of hfe, hfev- and quartz. In Yvo Desmedt, editor, *Public Key Cryptography - PKC 2003, 6th International Workshop on Theory and Practice in Public Key Cryptography, Miami, FL, USA, January 6-8, 2003, Proceedings*, volume 2567 of *Lecture Notes in Computer Science*, pages 337–350. Springer, 2003.

[14] Nicolas Courtois, Matthieu Finiasz, and Nicolas Sendrier. How to achieve a mceliece-based digital signature scheme. In Colin Boyd, editor, *Advances in Cryptology - ASIACRYPT 2001, 7th International Conference on the Theory and Application of Cryptology and Information Security, Gold Coast, Australia, December 9-13, 2001, Proceedings*, volume 2248 of *Lecture Notes in Computer Science*, pages 157–174. Springer, 2001.

[15] Jintai Ding. New cryptographic constructions using generalized learning with errors problem. *IACR Cryptology ePrint Archive*, 2012:387, 2012.

[16] Léo Ducas. Accelerating bliss: the geometry of ternary polynomials. *IACR Cryptology ePrint Archive*, 2014:874, 2014.

[17] Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 40–56. Springer, 2013.

[18] Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.

[19] Léo Ducas and Phong Q. Nguyen. Learning a zonotope and more: Cryptanalysis of ntrusign countermeasures. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, volume 7658 of *Lecture Notes in Computer Science*, pages 433–450. Springer, 2012.

[20] L. Eldar and P. W. Shor. An Efficient Quantum Algorithm for a Variant of the Closest Lattice-Vector Problem. *ArXiv e-prints*, nov 2016.

[21] Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Side-channel attacks on BLISS lattice-based signatures: Exploiting branch tracing against strongswan and electromagnetic emanations in microcontrollers. In Thuraisingham et al. [48], pages 1857–1874.

[22] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(133), 2007.

[23] Lov K. Grover. A fast quantum mechanical algorithm for database search. In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 212–219. ACM, 1996.

[24] Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. NTRUSIGN: digital signatures using the NTRU lattice. In Marc Joye, editor, *Topics in Cryptology - CT-RSA 2003, The Cryptographers' Track at the RSA Conference 2003, San Francisco, CA, USA, April 13-17, 2003, Proceedings*, volume 2612 of *Lecture Notes in Computer Science*, pages 122–140. Springer, 2003.

[25] Andreas Hülsing, Lea Rausch, and Johannes Buchmann. Optimal parameters for XMSS MT. *IACR Cryptology ePrint Archive*, 2017:966, 2017.

[26] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In Bo-Yin Yang, editor, *Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011, Taipei, Taiwan, November 29 - December 2, 2011. Proceedings*, volume 7071 of *Lecture Notes in Computer Science*, pages 19–34. Springer, 2011.

[27] Marc Kaplan, Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 207–237. Springer, 2016.

[28] Aviad Kipnis, Jacques Patarin, and Louis Goubin. Unbalanced oil and vinegar signature schemes. In Jacques Stern, editor, *Advances in Cryptology - EUROCRYPT '99, International Conference on the Theory and Application of Cryptographic Techniques, Prague, Czech Republic, May 2-6, 1999, Proceeding*, volume 1592 of *Lecture Notes in Computer Science*, pages 206–222. Springer, 1999.

[29] Aviad Kipnis and Adi Shamir. Cryptanalysis of the oil & vinegar signature scheme. In Hugo Krawczyk, editor, *Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference, Santa Barbara, California, USA, August 23-27, 1998, Proceedings*, volume 1462 of *Lecture Notes in Computer Science*, pages 257–266. Springer, 1998.

[30] Paul Kirchner and Pierre-Alain Fouque. Comparison between subfield and straightforward attacks on NTRU. *IACR Cryptology ePrint Archive*, 2016:717, 2016.

[31] Kazukuni Kobara and Hideki Imai. Semantically secure mceliece public-key cryptosystems-conversions for mceliece PKC. In Kwangjo Kim, editor, *Public Key Cryptography, 4th International Workshop on Practice and Theory in Public Key Cryptography, PKC 2001, Cheju Island, Korea, February 13-15, 2001, Proceedings*, volume 1992 of *Lecture Notes in Computer Science*, pages 19–35. Springer, 2001.

[32] Pascal Koiran, Vincent Nesme, and Natacha Portier. A quantum lower bound for the query complexity of simon's problem. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11-15, 2005, Proceedings*, volume 3580 of *Lecture Notes in Computer Science*, pages 1287–1298. Springer, 2005.

[33] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptography*, 75(3):565–599, 2015.

[34] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-lwe cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 35–54. Springer, 2013.

[35] Tsutomu Matsumoto and Hideki Imai. Public quadratic polynominal-tuples for efficient signature-verification and message-encryption. In Christoph G. Günther, editor, *Advances in Cryptology - EUROCRYPT '88, Workshop on the Theory and Application of of Cryptographic Techniques, Davos, Switzerland, May 25-27, 1988, Proceedings*, volume 330 of *Lecture Notes in Computer Science*, pages 419–453. Springer, 1988.

[36] R. J. McEliece. A Public-Key Cryptosystem Based On Algebraic Coding Theory. *Deep Space Network Progress Report*, 44:114–116, jan 1978.

[37] Ralph C. Merkle. A certified digital signature. In Gilles Brassard, editor, *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 218–238. Springer, 1989.

[38] Rafael Misoczki, Jean-Pierre Tillich, Nicolas Sendrier, and Paulo S. L. M. Barreto. Mdpc-mceliece: New mceliece variants from moderate density parity-check codes. In *Proceedings of the 2013*

IEEE International Symposium on Information Theory, Istanbul, Turkey, July 7-12, 2013, pages 2069–2073. IEEE, 2013.

[39] Phong Q. Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *J. Cryptology*, 22(2):139–160, 2009.

[40] Harald Niederreiter. Knapsack type cryptosystems and algebraic coding theory. 15, 01 1986.

[41] Jacques Patarin. Cryptanalysis of the matsumoto and imai public key scheme of eurocrypt'88. In Don Coppersmith, editor, *Advances in Cryptology - CRYPTO '95, 15th Annual International Cryptology Conference, Santa Barbara, California, USA, August 27-31, 1995, Proceedings*, volume 963 of *Lecture Notes in Computer Science*, pages 248–261. Springer, 1995.

[42] Jacques Patarin. Hidden fields equations (HFE) and isomorphisms of polynomials (IP): two new families of asymmetric algorithms. In Ueli M. Maurer, editor, *Advances in Cryptology - EUROCRYPT '96, International Conference on the Theory and Application of Cryptographic Techniques, Saragossa, Spain, May 12-16, 1996, Proceeding*, volume 1070 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 1996.

[43] Peter Pessl, Leon Groot Bruinderink, and Yuval Yarom. To BLISS-B or not to be: Attacking strongswan's implementation of post-quantum signatures. In Thuraisingham et al. [48], pages 1843–1855.

[44] Albrecht Petzoldt, Ming-Shing Chen, Bo-Yin Yang, Chengdong Tao, and Jintai Ding. Design principles for hfev- based multivariate signature schemes. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part I*, volume 9452 of *Lecture Notes in Computer Science*, pages 311–334. Springer, 2015.

[45] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93. ACM, 2005.

[46] Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *35th Annual Symposium on Foundations of Computer Science, Santa Fe, New Mexico, USA, 20-22 November 1994*, pages 124–134. IEEE Computer Society, 1994.

[47] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, volume 6632 of *Lecture Notes in Computer Science*, pages 27–47. Springer, 2011.

[48] Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*. ACM, 2017.

We offer Denmark's strongest and most versatile competences within consultancy, development and application of information technologies

ALEXANDRA INSTITUTE