

Post-Quantum Key Exchange for the Internet and the Open Quantum Safe Project*

Douglas Stebila^{1,†} and Michele Mosca^{2,3,4,‡}

¹ *Department of Computing and Software,
McMaster University, Hamilton, Ontario, Canada*
stebilad@mcmaster.ca

² *Institute for Quantum Computing and Department of Combinatorics & Optimization,
University of Waterloo, Waterloo, Ontario, Canada*

³ *Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada*

⁴ *Canadian Institute for Advanced Research, Toronto, Ontario, Canada*
mmosca@uwaterloo.ca

July 28, 2017

Abstract

Designing public key cryptosystems that resist attacks by quantum computers is an important area of current cryptographic research and standardization. To retain confidentiality of today's communications against future quantum computers, applications and protocols must begin exploring the use of quantum-resistant key exchange and encryption. In this paper, we explore post-quantum cryptography in general and key exchange specifically. We review two protocols for quantum-resistant key exchange based on lattice problems: BCNS15, based on the ring learning with errors problem, and Frodo, based on the learning with errors problem. We discuss their security and performance characteristics, both on their own and in the context of the Transport Layer Security (TLS) protocol. We introduce the Open Quantum Safe project, an open-source software project for prototyping quantum-resistant cryptography, which includes liboqs, a C library of quantum-resistant algorithms, and our integrations of liboqs into popular open-source applications and protocols, including the widely used OpenSSL library.

*Based on the Stafford Tavares Invited Lecture at Selected Areas in Cryptography (SAC) 2016 by D. Stebila.

†Supported in part by Australian Research Council (ARC) Discovery Project grant DP130104304, Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grant RGPIN-2016-05146, and an NSERC Discovery Accelerator Supplement grant.

‡Supported by NSERC, CFI, and ORF. IQC and the Perimeter Institute are supported in part by the Government of Canada and the Province of Ontario.

Contents

1	Introduction	3
2	Lattice-based cryptography and the LWE problems	5
2.1	The Learning with Errors problem	5
2.2	The Ring Learning with Errors problem	7
3	Key exchange protocols from LWE and ring-LWE	8
3.1	Common tools: reconciliation	8
3.2	Ring-LWE-based key exchange: BCNS15	9
3.3	LWE-based key exchange: Frodo	10
3.4	Performance of post-quantum key exchange	11
3.5	From unauthenticated to authenticated key exchange	13
4	Integrating post-quantum key exchange into TLS	13
4.1	Performance of post-quantum key exchange in TLS	13
5	Interlude: programming is hard	15
6	Open Quantum Safe: a software framework for post-quantum cryptography	15
6.1	liboqs	16
6.2	Application/protocol integrations	17
6.3	Case study: adding NewHope to liboqs and OpenSSL	18
7	Conclusion and outlook	18
8	Acknowledgements	19
	References	19

1 Introduction

All Internet security protocols that use cryptography, such as the Transport Layer Security (TLS, a.k.a. the Secure Sockets Layer (SSL)) protocol [18] have the same basic structure: *public key cryptography* is used to authenticate the communicating parties to each other and to establish a shared secret key, which is then used in *symmetric cryptography* to provide confidentiality and integrity to their communication. The security of most public key cryptosystems depends on the difficulty of solving some mathematical problem, such as factoring large numbers or computing discrete logarithms in finite field or elliptic curve groups. The best known solutions to these problems run in exponential (or sub-exponential) time, making it infeasible for attackers to break the schemes.

Quantum mechanics allows for devices that operate on quantum bits, known as *qubits*, which are two-state quantum systems that can be in any quantum superposition of 0 and 1. Such devices are called quantum computers, and could solve certain types of problems much faster than “classical” (non-quantum) computers. Shor’s algorithm [48] could efficiently (i.e., in polynomial time) factor large numbers and compute discrete logarithms, breaking all widely deployed public key cryptosystems. Most symmetric key schemes, such as the Advanced Encryption Standard (AES) cipher, would not be broken by quantum algorithms, although would generally need bigger keys. While large-scale quantum computers do not yet exist, building quantum computers is an active area of research. and Schoelkopf [17] identify seven stages in the development of quantum computers: so far, physicists can perform operations on single and multiple physical qubits, perform non-destructive measurements for error correction, and are making progress on constructing logical memories with longer lifetime than physical qubits; to achieve large-scale quantum computation, we will require the ability to perform operations on single and multiple logical qubits with fault-tolerant computation. Regarding the million-dollar question of when a large-scale quantum computer will be built, in 2015 Mosca [38] stated “I estimate a 1/7 chance of breaking RSA-2048 by 2026 and a 1/2 chance by 2031.”

Any attacker who records present-day communications would be able to decrypt it once a quantum computer is built; and there is evidence that governments are storing vast quantities of encrypted Internet traffic. This motivates the urgent use of cryptography that is designed to be safe against quantum attackers—called “post-quantum” or “quantum-safe” or “quantum-resistant” cryptography. In August 2015, the United States National Security Agency (NSA) issued a memo regarding its Suite B cryptographic algorithms for government use, advising that it plans to “transition to quantum resistant algorithms in the not too distant future” [39]. In August 2016, the United States National Institute of Standards and Technology (NIST) launched its post-quantum crypto project¹, a multi-year process with the goal of evaluating and standardizing one or more quantum-resistant public key cryptosystems.

Post-quantum cryptography. There are several classes of mathematical problems that are conjectured to resist attacks by quantum computers and have been used to construct public key cryptosystems, several of which date from the early days of public key cryptography. These include:

- *Code-based cryptography.* The McEliece public key encryption scheme [36] was one of the first public key schemes, and is based on error-correcting codes, in particular, the hardness of decoding a general linear code. Niederreiter [40] subsequently proposed a digital signature scheme based on error correcting codes.
- *Hash-based cryptography.* Merkle [37] first proposed the use of hash functions for digitally signing documents; Lamport and Diffie [30] and Winternitz then showed how to convert

¹<http://www.nist.gov/pqcrypto>

Merkle’s one-time signature scheme into a many-time signature scheme. These schemes are based entirely on standard hash function properties, and thus are believed to be among the most quantum-resistant. Modern variants include SPHINCS [7] and XMSS [13].

- *Multivariate cryptography.* These cryptosystems are based on the difficulty of solving non-linear, usually quadratic, polynomials, over a field [41, 35].
- *Lattice-based cryptography.* Ajtai [1] proposed the first cryptographic schemes directly based on lattices. Regev [46] then introduced the related *learning with errors (LWE) problem*, the security of which is based on lattice problems, and which now forms the basis of a variety of public key encryption and signature schemes [31]. The *ring learning with errors (ring-LWE) problem* [33] uses additional structure which allows for smaller key sizes. Another scheme whose security relates to lattices is the NTRU scheme [26], which also allows for fairly small key sizes.
- *Supersingular elliptic curve isogenies.* One of the newest candidates for quantum-resistant public key cryptography is based on the difficulty of finding isogenies between supersingular elliptic curves [20].

In addition, quantum information can be used directly to create cryptosystems; this is called *quantum cryptography*. For example, quantum key distribution allows two parties to establish a shared secret key using quantum communication and an authenticated classical channel. While this can provide very strong security, it is not yet a candidate for widespread usage since it requires physical infrastructure capable of transmitting quantum states reliably over long distances, so in the rest of this paper we focus solely on quantum-resistant cryptography using classical (non-quantum) computers.

Existing quantum-resistant schemes generally have several limitations. Compared with traditional RSA, finite field, and elliptic curve discrete logarithm schemes, all quantum-resistant schemes have either larger public keys, larger ciphertexts/signatures, or slower runtime. Many quantum-resistant schemes are also based on mathematical problems that are, from a cryptographic perspective, quite new, and thus have received comparably less cryptanalysis. There remain many open questions in post-quantum cryptography, making it an exciting and active research area: the design of better public key encryption and signature schemes with smaller keys and ciphertexts/signatures; improved cryptanalysis leading to better parameter estimates; development of fast, secure implementations suitable for high-performance servers and small embedded devices; and integration into existing network infrastructure and applications.

(It is worth noting that research into post-quantum cryptography is valuable even if large-scale quantum computers are never built: it is possible that the factoring, RSA, or discrete logarithms problems will be solved by some (non-quantum) mathematical breakthrough. Having a diverse family of cryptography assumptions on which we can base public key cryptography protects against such a scenario. Furthermore, the cryptographic agility that will help prepare for a transition to yet-to-be-determined quantum-resistant cryptographic algorithms will also enable the ability to respond quickly to other unexpected weaknesses in cryptographic algorithms.)

This paper. In this paper, we discuss two research projects in the area of lattice-based key exchange: the “BCNS15” protocol [10] based on the ring-LWE problem, and the “Frodo” protocol [9] based on the LWE problem. We will explain the basic mathematics of these protocols, and our results on the performance of these protocols and their integration into the TLS protocol. We will introduce the Open Quantum Safe project, an open-source software project designed for evaluating post-quantum cryptography candidates and prototyping their use in applications and protocols such as TLS.

This line of work focuses initially on key exchange, with digital signatures to follow closely. As noted above, any attacker who records present-day communications protected using non-quantum-resistant cryptography would be able to decrypt it once a quantum computer is built. This implies that information that needs to remain confidential for many years needs to be protected with quantum-resistant cryptography even before quantum computers exist. In communication protocols like TLS, digital signatures are used to authenticate the parties and key exchange is used to establish a shared secret, which can then be used in symmetric cryptography. This means that, for security against a future quantum adversary, authentication in today’s secure channel establishment protocols can still rely on traditional primitives (such as RSA or elliptic curve signatures), but we should incorporate post-quantum key exchange to provide quantum-resistant long-term confidentiality. This has the benefit of allowing us to introduce new post-quantum ciphersuites in TLS while relying on the existing RSA-based public key infrastructure for certificate authorities. However, applications which require long-term integrity, such as signing contracts and archiving documents, will need to begin considering quantum-resistant signature schemes.

Notation. Let χ be a distribution; $a \stackrel{\$}{\leftarrow} \chi$ denotes sampling a randomly according to χ . The uniform distribution is denoted by \mathcal{U} . Vectors are denoted in lower-case bold, like \mathbf{a} ; matrices are denoted in upper-case bold, like \mathbf{A} . The inner product of two vectors \mathbf{a} and \mathbf{b} is $\langle \mathbf{a}, \mathbf{b} \rangle$. Sampling each component of the length- n vector \mathbf{a} independently at random from χ is denoted by $\mathbf{a} \stackrel{\$}{\leftarrow} \chi^n$. If \mathcal{A} is a probabilistic algorithm, then $y \stackrel{\$}{\leftarrow} \mathcal{A}(x)$ denotes running \mathcal{A} on input x with fresh randomness and storing the output in variable y , and $y \stackrel{\$}{\leftarrow} \mathcal{A}^O(x)$ denotes running \mathcal{A} with oracle access to procedure O .

2 Lattice-based cryptography and the LWE problems

In a seminal 1996 work, Ajtai [1] first proposed a cryptographic construction (in that case, a hash function) that relied on the hardness of a computational problem on lattices (the Short Integer Solution (SIS) problem). A subsequent work by Ajtai and Dwork [2] presented a public key encryption scheme based on another lattice problem. Concurrently, Hoffstein, Pipher, and Silverman [26] created the NTRU public key encryption scheme which can be viewed as involving algebraically structured lattices. A variety of research on the use of lattices in constructing cryptosystems continued during that era, and a detailed chronology is outside the scope of this paper; see one of the many surveys of lattice-based cryptography, such as Peikert’s [44].

2.1 The Learning with Errors problem

In 2005, Regev [46] introduced the *learning with errors (LWE) problem*, showed that LWE is related to the hardness of a lattice problem (the Gap Shortest Vector Problem (GapSVP)), and gave a public key encryption scheme based on LWE. Being a more abstract algebraic problem, LWE can be easier to work with in terms of building cryptosystems, and a large amount of research into the hardness of LWE and its use in cryptography has followed; again, see a survey such as [44] for a detailed chronology.

The search learning with errors problem is like a noisy version of solving a system of linear equations: given a matrix \mathbf{A} and a vector $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$, find \mathbf{s} .

Definition 1 (Search LWE problem). Let n , m , and q be positive integers. Let χ_s and χ_e be distributions over \mathbb{Z} . Let $\mathbf{s} \stackrel{\$}{\leftarrow} \chi_s^n$. Let $\mathbf{a}_i \stackrel{\$}{\leftarrow} \mathcal{U}(\mathbb{Z}_q^n)$, $e_i \stackrel{\$}{\leftarrow} \chi_e$, and set $b_i \leftarrow \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \pmod q$, for

$i = 1, \dots, m$. The *search LWE problem* for $(n, m, q, \chi_s, \chi_e)$ is to find \mathbf{s} given $(\mathbf{a}_i, b_i)_{i=1}^m$. In particular, for algorithm \mathcal{A} , define the advantage

$$\text{Adv}_{n,m,q,\chi_s,\chi_e}^{\text{lwe}}(\mathcal{A}) = \Pr \left[\mathbf{s} \xleftarrow{\$} \chi_s^n; \mathbf{a}_i \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^n); e_i \xleftarrow{\$} \chi_e; \right. \\ \left. b_i \leftarrow \langle \mathbf{a}_i, \mathbf{s}_i \rangle + e \bmod q : \mathcal{A}((\mathbf{a}_i, b_i)_{i=1}^m) = \mathbf{s} \right] .$$

For appropriate distributions χ_s and χ_e , not only is it conjectured to be hard to find the secret vector \mathbf{s} , it is even conjectured that LWE samples $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e)$ look independent and random: this is the decision LWE problem.

Definition 2 (Decision LWE problem). Let n and q be positive integers. Let χ_s and χ_e be distributions over \mathbb{Z} . Let $\mathbf{s} \xleftarrow{\$} \chi_s^n$. Define the following two oracles:

- $O_{\chi_e, \mathbf{s}}$: $\mathbf{a} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^n)$, $e \xleftarrow{\$} \chi_e$; return $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e \bmod q)$.
- U : $\mathbf{a} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^n)$, $u \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q)$; return (\mathbf{a}, u) .

The *decision LWE problem* for (n, q, χ_s, χ_e) is to distinguish $O_{\chi_s, \mathbf{s}}$ from U . In particular, for algorithm \mathcal{A} , define the advantage

$$\text{Adv}_{n,q,\chi_s,\chi_e}^{\text{dlwe}}(\mathcal{A}) = \left| \Pr(\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n : \mathcal{A}^{O_{\chi_e, \mathbf{s}}}() = 1) - \Pr(\mathcal{A}^U() = 1) \right| .$$

Choice of distributions. The error distribution χ_e is usually a discrete Gaussian distribution of width αq for “error rate” $\alpha < 1$.

LWE was originally phrased involving a uniform distribution on the secret \mathbf{s} ($\chi_s^n = \mathcal{U}(\mathbb{Z}_q^n)$). Applebaum et al. [5] showed that the *short secrets* (or “normal form”) variant, in which $\chi_s = \chi_e$, has a tight reduction to the original uniform secrets variant. In what follows, we use the short secrets variant throughout, and abbreviate to a single error distribution χ (using shorthand notation $\text{Adv}_{n,m,q,\chi}^{\text{lwe}}$ and $\text{Adv}_{n,q,\chi}^{\text{dlwe}}$).

Difficulty. Difficulty of both search and decision LWE problems depends on the size of n , m , and q , as well as the distributions χ_s and χ_e . Regev [46] showed that, for appropriate parameters, search and decision LWE are worst-case hard assuming the (average case) hardness of a lattice problem. In particular, he showed first that search-LWE is at least as hard as solving the worst-case lattice problems GapSVP_γ and SIVP_γ (for a parameter γ depending on n and α) using a quantum reduction; then that decision-LWE is at least as hard as the search version using a classical reduction. A sequence of later results have improved various aspects (making the first reduction classical, not quantum; handling more moduli); see Peikert’s survey [44, §4.2.2] for a list.

Extracting secret bits. The decision LWE problem effectively yields an element $\langle \mathbf{a}, \mathbf{s} \rangle + e \in \mathbb{Z}_q$ that is indistinguishable from random. Parties using LWE to establish a shared secret for public key encryption (like in Regev’s scheme) or key agreement (as we will see in the next section) will only approximately agree on the same value modulo q , so they will have to apply some reconciliation function and extracting a small number of bits (maybe even just 1 bit) from a single element of \mathbb{Z}_q . In order to establish a multi-bit shared secret with LWE, the parties will hence need to send many samples, which we can then think of in matrix form: a matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{m \times n}$, and an error $\mathbf{e} \xleftarrow{\$} \chi^n$, to obtain $\mathbf{b} \leftarrow \mathbf{A}\mathbf{s} + \mathbf{e} \in \mathbb{Z}_q^m$. This increases communication sizes m -fold, and requires approximately $O(mn \log q)$ bits of communication to obtain an m -bit secret. To reduce communication sizes, one could try to introduce some structure to the matrix \mathbf{A} , for example making each row the cyclic shift of the previous row. However, rather than working in matrix form, we can shift our representation to a polynomial ring, leading us to the ring-LWE problem.

2.2 The Ring Learning with Errors problem

In 2010, Lyubashevsky, Peikert, and Regev [34] introduced the ring-LWE problem. Let $R = \mathbb{Z}[X]/\langle X^n + 1 \rangle$, where n is a power of 2. Let q be an integer, and define $R_q = R/qR$, i.e., $R_q = \mathbb{Z}_q[X]/\langle X^n + 1 \rangle$. In other words, R_q consists of polynomials of degree at most $n - 1$, with coefficients in \mathbb{Z}_q , and the wrapping rule that $X^n \equiv -1 \pmod{q}$. The search and decision ring-LWE problems are analogues of the corresponding LWE problems, except with ring elements rather than vectors.

Definition 3 (Search ring-LWE problem). Let n and q be positive integers. Let χ_s and χ_e be distributions over R_q . Let $s \xleftarrow{\$} \chi_s$. Let $a \xleftarrow{\$} \mathcal{U}(R_q)$, $e \xleftarrow{\$} \chi_e$, and set $b \leftarrow as + e$. The *search ring-LWE problem* for (n, q, χ_s, χ_e) is to find s given (a, b) . In particular, for algorithm \mathcal{A} define the advantage

$$\text{Adv}_{n,q,\chi_s,\chi_e}^{\text{rlwe}}(\mathcal{A}) = \Pr [s \xleftarrow{\$} \chi_s; a \xleftarrow{\$} \mathcal{U}(R_q); e \xleftarrow{\$} \chi_e; b \leftarrow as + e : \mathcal{A}(a, b) = s] .$$

Again, for appropriate distributions χ_s and χ_e , not only is it conjectured to be hard to find the secret s , it is even conjectured that ring-LWE samples $(a, as + e)$ look independent and random: this is the decision ring-LWE problem.

Definition 4 (Decision ring-LWE problem). Let n and q be positive integers. Let χ_s and χ_e be distributions over R_q . Let $s \xleftarrow{\$} \chi_s$. Define the following two oracles:

- $O_{\chi_e,s}$: $a \xleftarrow{\$} \mathcal{U}(R_q)$, $e \xleftarrow{\$} \chi_e$; return $(a, as + e)$.
- U : $a, u \xleftarrow{\$} \mathcal{U}(R_q)$; return (a, u) .

The *decision ring-LWE problem* for (n, q, χ_s, χ_e) is to distinguish $O_{\chi_e,s}$ from U . In particular, for algorithm \mathcal{A} , define the advantage

$$\text{Adv}_{n,q,\chi_s,\chi_e}^{\text{drlwe}}(\mathcal{A}) = \left| \Pr(s \xleftarrow{\$} R_q : \mathcal{A}^{O_{\chi_e,s}}() = 1) - \Pr(\mathcal{A}^U() = 1) \right| .$$

Choice of distributions. The error distribution χ_e is usually a discretized Gaussian distribution in the canonical embedding of R ; for an appropriate choice of parameters, we can sample ring elements from χ_e by sampling each coefficient of the polynomial independently from a related distribution.

As with LWE, ring-LWE can be formulated using either a uniform secret ($\chi_s = \mathcal{U}(R_q)$) or with short secrets ($\chi_s = \chi_e$), which has a tight reduction to the original secrets variant. In what follows, we use the short secrets variant throughout, and abbreviate to a single error distribution χ (using shorthand notation $\text{Adv}_{n,q,\chi}^{\text{rlwe}}$ and $\text{Adv}_{n,q,\chi}^{\text{drlwe}}$).

Difficulty. Difficulty of both search and decision ring-LWE depends on the parameters n and q and the distributions χ_s and χ_e . Lyubashevsky et al. [34] showed that search ring-LWE as hard as quantumly solving approximate shortest vector problem on an *ideal* lattice in R ; and then the classical search-to-decision reduction applies.

Because of the additional structure present in ring-LWE, the choice of n and q requires greater care than the unstructured LWE problem [45]. There is also the risk that the ring-LWE problem may be easier than the LWE problem. Currently, the best known algorithms for solving hard problems in ideal lattices [14, 29] are the same as those for regular lattices (ignoring small polynomial speedups); and in some sieving algorithms, the ideal case enables one to save a small constant factor of time or space [47, 11]. Very recently Cramer et al. [16] gave a quantum polynomial time algorithm

algorithm for ideal-SVP with certain parameters, but this is not currently applicable to ring-LWE. In summary, some view LWE as a more “conservative” security choice than ring-LWE, though there is no appreciable security difference at present.

Extracting secret bits. The decision ring-LWE problem effectively yields a ring element that is indistinguishable from random. Being an element of $R_q = \mathbb{Z}_q[X]/\langle X^n + 1 \rangle$, we have n coefficients each of which is an element of \mathbb{Z}_q . As with LWE, cryptographic constructions using this will need to reconcile approximately equal shared secrets, and thus can extract only a small number of bits (maybe even just 1 bit) from each coefficient. But since there are n (independent-looking) coefficients, one can extract n random-looking bits from a single ring element. Thus, one needs approximately $O(n \log q)$ bits of communication to obtain an n -bit secret, a substantial reduction compared to LWE. Thus, in practice, one must decide between the decreased communication of ring-LWE versus the potentially more conservative security of LWE.

3 Key exchange protocols from LWE and ring-LWE

Regev [46] was the first to give a public key encryption scheme from the learning with errors problem, and Lyubashevsky, Peikert, and Regev [33] were the first to give a public key encryption scheme from ring-LWE. Like ElGamal public key encryption, both these schemes implicitly contain a key encapsulation mechanism and then one-time-mask the KEM shared secret with (an encoded form of) the message. Peikert [42] describes a corresponding approximate LWE key agreement protocol. In 2010, Lindner and Peikert [31] gave an improved LWE-based public key encryption scheme, and a ring-LWE analogue, and described how to view it as an approximate key agreement protocol. This was followed by detailed LWE- and ring-LWE-based key agreement protocols by Ding et al. [19] (including a single-bit reconciliation mechanism to obtain exact key agreement); a sketch of an LWE-based key agreement scheme by Blazy et al. [8, Fig. 1, 2]; and detailed ring-LWE-based key encapsulation mechanisms by Fujioaka et al. [22, §5.2] and Peikert [43] (with an alternative single-bit reconciliation mechanism). In addition to basic unauthenticated key exchange, there have been works on using LWE to create *password-authenticated* key exchange [28] and using ring-LWE to create *authenticated* key exchange [49] (though the security proof of the latter is questioned [24]).

In this section, we will examine two unauthenticated key agreement protocols in which this paper’s first author was involved. Frodo [9], an LWE-based key exchange protocol, is an instantiation of the Lindner–Peikert LWE approximate key agreement scheme using a generalization of Peikert’s reconciliation mechanism in which multiple bits are extracted from a single element of \mathbb{Z}_q . BCNS15 [10], a ring-LWE-based key exchange protocol, is an instantiation of the key exchange scheme corresponding to the KEM in the Lyubashevsky–Peikert–Regev public key encryption scheme from ring-LWE using Peikert’s reconciliation mechanism.

3.1 Common tools: reconciliation

In both Frodo and BCNS15, the parties will establish an approximately equal shared secret, then exchange some “hints” that allow them to perform a reconciliation operation on the approximately equal shared secret to extract some secret bits that are, with high probability, the same for both parties. The reconciliation technique of Ding et al. [19] sends a single bit “hint” for each key bit and relies on the low-order bits of the shared secret; Peikert’s technique [43] also sends a single bit hint but relies on the high-order bits of the shared secret. The explanation below generalizes Peikert’s approach [43] to extract multiple bits.

Let $B \in \mathbb{N}$ be the number of bits we aim to extract from one element of \mathbb{Z}_q . Assume $B < (\log_2 q) - 1$. Let $\bar{B} = \lceil \log_2 q \rceil - B$. Let $v \in \mathbb{Z}_q$, represented canonically as an integer in $[0, q)$. Define the *rounding* function

$$\lfloor \cdot \rfloor_{2^B} : \mathbb{Z}_q \rightarrow \mathbb{Z}_{2^B} : v \mapsto \left\lfloor 2^{-\bar{B}} v \right\rfloor \bmod 2^B ,$$

where $\lfloor \cdot \rfloor : \mathbb{R} \rightarrow \mathbb{Z}$ rounds real number x to the closest integer. When q is a multiple of 2^B , $\lfloor \cdot \rfloor_{2^B}$ outputs the B most significant bits of $(v + 2^{\bar{B}-1}) \bmod q$, thereby partitioning \mathbb{Z}_q into 2^B intervals of integers with the same B most significant bits (up to a cyclic shift of the values that centres these intervals around 0).

Define the *cross-rounding* function

$$\langle \cdot \rangle_{2^B} : \mathbb{Z}_q \rightarrow \mathbb{Z}_2 : v \mapsto \left\lfloor 2^{-\bar{B}+1} v \right\rfloor \bmod 2 ,$$

where $\lfloor \cdot \rfloor : \mathbb{R} \rightarrow \mathbb{Z}$ takes the floor of the real number x . When q is a multiple of 2^{B+1} , $\langle \cdot \rangle_{2^B}$ partitions \mathbb{Z}_q into two intervals based according to their $(B+1)$ th most significant bit.

On input of $w \in \mathbb{Z}_q$ and $c \in \{0, 1\}$, the *reconciliation* function $\text{rec}_{2^B}(w, c)$ outputs $\lfloor v \rfloor_{2^B}$, where v is the *closest* element to w such that $\langle v \rangle_{2^B} = c$.

If Alice and Bob have approximately equal values $v, w \in \mathbb{Z}_q$, they can use the following process to derive B bits that are, with high probability, equal. Suppose q is a multiple of 2^B . Bob computes $c \leftarrow \langle v \rangle_{2^B}$ and sends c to Alice. Bob computes $k' \leftarrow \lfloor v \rfloor_{2^B}$. Alice computes $k \leftarrow \text{rec}_{2^B}(w, c)$.

Security of this technique follows from the following fact: if $v \in \mathbb{Z}_q$ is uniformly random, then $\lfloor v \rfloor_{2^B}$ is uniformly random given $\langle v \rangle_{2^B}$.

Correctness follows if v and w are sufficiently close. Namely, if $|v - w \pmod{q}| < 2^{\bar{B}-2}$, then $\text{rec}_{2^B}(w, \langle v \rangle_{2^B}) = \lfloor v \rfloor_{2^B}$. Parameters must be chosen so that v and w are sufficiently close.

For our parameters in the ring setting, we will want to extract 1 bit from each element of \mathbb{Z}_q , but q will not be a multiple of 2. Peikert suggested the following technique: Bob computes $\bar{v} \stackrel{\$}{\leftarrow} \text{dbl}(v)$, where $\text{dbl} : \mathbb{Z}_q \rightarrow \mathbb{Z}_{2q} : x \mapsto 2x - e$, where e is sampled from $\{-1, 0, 1\}$ with probabilities $p_{-1} = p_1 = \frac{1}{4}$ and $p_0 = \frac{1}{2}$. Bob computes $c \leftarrow \langle \bar{v}/2 \rangle_2$ and sends c to Alice. Bob computes $k' \leftarrow \lfloor \bar{v}/2 \rfloor_2$. Alice computes $k \leftarrow \text{rec}_2(2w, c)$.

For ring-LWE, these functions are extended from \mathbb{Z}_q to the ring $R_q = \mathbb{Z}_q[X]/\langle X^n + 1 \rangle$ coefficient-wise. For matrix forms of LWE, these functions can be extended to vectors component-wise.

3.2 Ring-LWE-based key exchange: BCNS15

Protocol. The BCNS15 protocol [10], based on the ring-LWE problem, is shown in Figure 1. Alice and Bob exchange ring-LWE samples $b = as + e$ and $b' = as' + e'$. They can then compute an approximately equal shared secret:

$$sb' = sas' + se' \approx sas' + s'e = bs' \in R_q = \mathbb{Z}_q[X]/\langle X^n + 1 \rangle .$$

From each coefficient of the approximately equal shared secret, they extract a single secret bit.

Security. Assuming the decision ring-LWE problem is hard for the parameters chosen, the BCNS15 key exchange protocol is a secure unauthenticated key exchange protocol. The argument follows [31, 43] by using two applications of the decision ring-LWE assumption: first, on Alice's computations involving s (so b becomes independent from s), and second on Bob's computations involving s' (so b' and v become independent from s'). This makes the approximately equal shared secret v uniformly random from the adversary's perspective, and as noted above the hint c reveals no information about extracted key k' .

Public parameters	
Decision ring-LWE parameters n, q, χ	
$a \xleftarrow{\$} \mathcal{U}(R_q)$	
Alice	Bob
$s, e \xleftarrow{\$} \chi$	
$b \leftarrow as + e \in R_q$	\xrightarrow{b}
	$s', e' \xleftarrow{\$} \chi$
	$b' \leftarrow as' + e' \in R_q$
	$e'' \xleftarrow{\$} \chi$
	$v \leftarrow bs' + e'' \in R_q$
	$\bar{v} \xleftarrow{\$} \text{dbl}(v) \in R_{2q}$
	$\xleftarrow{b', c}$
$k_A \leftarrow \text{rec}_2(2b's, c) \in \{0, 1\}^n$	$c \leftarrow \langle \bar{v}/2 \rangle_2 \in \{0, 1\}^n$
	$k_B \leftarrow \lfloor \bar{v}/2 \rfloor_2 \in \{0, 1\}^n$

Figure 1: **BCNS15: Unauthenticated Diffie–Hellman-like key exchange from ring-LWE**

Parameters. The BCNS15 protocol is instantiated with $n = 1024$ and $q = 2^{32} - 1$. The error distribution χ is a discrete Gaussian distribution; because n is a power of 2, this can be achieved by sampling each coefficient from a discrete Gaussian $D_{\mathbb{Z}, \sigma}$ with $D_{\mathbb{Z}, \sigma}(x) = \frac{1}{S} e^{-x^2/(2\sigma^2)}$ where $S = 1 + 2 \sum_{k=1}^{\infty} e^{-k^2/(2\sigma^2)}$. With these parameters, the probability that reconciliation yields $k \neq k'$ is much less than 2^{-128} . Total communication required for two parties to establish a shared secret is 8,320 bytes.

Based on hardness estimates by Albrecht et al. [3], breaking the system with these parameters would require $2^{163.8}$ operations on a classical computer with at least $2^{94.4}$ memory usage. Assuming a square-root speedup for quantum computers via Grover’s algorithm (though it is not known how to achieve a full square-root speedup), this suggests at least $2^{81.9}$ quantum security. Based on the same difficulty estimates for the subsequent NewHope protocol [4], Alkim et al. list BCNS15 as having 86-bit classical security and 78-bit quantum security.

Subsequent works. Alkim et al. [4] subsequently published the so-called “NewHope” protocol, making several improvements to the BCNS15 protocol. NewHope uses different parameters and a different error distribution (which was easier to sample), resulting in substantially improved performance and smaller communication (3,872 bytes). NewHope also uses a pseudorandomly generated a , rather than a fixed public parameter. In July 2016, Google announced that they were deploying a two-year experiment in the alpha version of their Chrome web browser (called “Canary”) that uses the NewHope key exchange protocol in a hybrid ciphersuite with elliptic curve Diffie–Hellman [12]. Further improvements to NewHope have been given by several papers [25, 32].

3.3 LWE-based key exchange: Frodo

Protocol. The Frodo key exchange protocol [9], based on the LWE problem, is shown in Figure 2. It uses a matrix form of the LWE problem: Alice uses m secrets $\mathbf{s}_1, \dots, \mathbf{s}_m$, represented as a matrix \mathbf{S} ; similarly for Bob. Alice and Bob exchange matrix LWE samples $\mathbf{B} = \mathbf{A}\mathbf{S} + \mathbf{E}$ and $\mathbf{B}' = \mathbf{S}'\mathbf{A}' + \mathbf{E}'$. They can then compute an approximately equal shared secret:

$$\mathbf{B}'\mathbf{S} = \mathbf{S}'\mathbf{A}\mathbf{S} + \mathbf{S}'\mathbf{E} \approx \mathbf{S}'\mathbf{A}\mathbf{S} + \mathbf{S}\mathbf{E}' = \mathbf{S}'\mathbf{B} \in \mathbb{Z}_q^{m \times m}.$$

Public parameters	
Decision LWE parameters n, q, χ ; integer m	
Alice	Bob
$seed \xleftarrow{\$} \{0, 1\}^\lambda$ $\mathbf{A} \leftarrow \text{PRF}(seed) \in \mathbb{Z}_q^{n \times n}$ $\mathbf{S}, \mathbf{E} \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times m})$ $\mathbf{B} \leftarrow \mathbf{AS} + \mathbf{E} \in \mathbb{Z}_q^{n \times m}$	$\mathbf{A} \leftarrow \text{PRF}(seed) \in \mathbb{Z}_q^{n \times n}$ $\mathbf{S}', \mathbf{E}' \xleftarrow{\$} \chi(\mathbb{Z}_q^{m \times n})$ $\mathbf{B}' \leftarrow \mathbf{S}'\mathbf{A} + \mathbf{E}' \in \mathbb{Z}_q^{m \times n}$ $\mathbf{E}'' \xleftarrow{\$} \chi(\mathbb{Z}_q^{m \times m})$ $\mathbf{V} \leftarrow \mathbf{S}'\mathbf{B} + \mathbf{E}'' \in \mathbb{Z}_q^{m \times m}$
$\mathbf{k} \leftarrow \text{rec}_{2^B}(\mathbf{B}'\mathbf{S}, \mathbf{C}) \in \mathbb{Z}_{2^B}^m$	$\mathbf{C} \leftarrow \langle \mathbf{V} \rangle_{2^B} \in \mathbb{Z}_{2^B}^{m \times m}$ $\mathbf{k}' \leftarrow \lfloor \mathbf{V} \rfloor_{2^B} \in \mathbb{Z}_{2^B}^m$

Figure 2: **Frodo: Unauthenticated Diffie–Hellman-like key exchange from LWE**

From each entry of the approximately equal shared secret, they extract B secret bits. Frodo follows NewHope’s idea of using a pseudorandomly generated \mathbf{A} .

Security. Assuming the decision LWE problem is hard for the parameters chosen, and PRF is a pseudorandom function, the Frodo key exchange protocol is a secure unauthenticated key exchange protocol. A hybrid argument goes from the standard decision-LWE problem to a matrix form of it, then the same argument as for BCNS15 above yields the indistinguishability of the session key.

Parameters. The Frodo paper contains several parameter sets, including a “recommended” parameter set, which uses $n = 752$, $q = 2^{15}$, $m = 8$, and $B = 4$. The error distribution χ is a concrete distribution specified in the paper, which is close in Renyi divergence to a rounded continuous Gaussian distribution (but requires fewer bits to sample). With these parameters, the probability that reconciliation yields $\mathbf{k} \neq \mathbf{k}'$ is $2^{-38.9}$. Total communication required for two parties to establish a shared secret is 8,320 bytes. The claimed security level is 140 bits of security against a classical adversary, and 130 bits against a quantum adversary. The paper also includes a higher-security “paranoid” parameter set, which conjectures a certain lower bound on lattice sieving for any adversary.

3.4 Performance of post-quantum key exchange

Table 1 (copied from [9]) shows the performance characteristics of several post-quantum key exchange protocols:

- BCNS ring-LWE key exchange, C implementation [10];
- NewHope ring-LWE key exchange, C implementation [4];
- NTRU public key encryption key transport using parameter set EES743EP1, C implementation;² and
- SIDH (supersingular isogeny Diffie–Hellman) key exchange, C implementation [15].

²<https://github.com/NTRUOpenSourceProject/ntru-crypto>

Scheme	Alice0	Bob	Alice1	Communication (bytes)		Claimed security	
	(ms)	(ms)	(ms)	A→B	B→A	classical	quantum
RSA 3072-bit	—	0.09	4.49	387 / 0*	384	128	—
ECDH <code>nistp256</code>	0.37	0.70	0.33	32	32	128	—
BCNS	1.01	1.59	0.17	4,096	4,224	86	78
NewHope	0.11	0.16	0.03	1,824	2,048	229	206
NTRU <code>EES743EP1</code>	2.00	0.28	0.15	1,027	1,022	256	128
Frodo Recomm.	1.13	1.34	0.13	11,377	11,296	144	130
Frodo Paranoid	1.25	1.64	0.15	13,057	12,976	177	161
SIDH	135	464	301	564	564	192	128

Table 1: **Performance of standalone cryptographic operations**, showing mean runtime in milliseconds of standalone cryptographic operations, communication sizes (public key / messages) in bytes, and claimed security level in bits. Table from [9]. * In TLS, the RSA public key is already included in the server’s certificate message, so RSA key transport imposes no additional communication from server to client.

The table also includes non-quantum-secure algorithms at the 128-bit classical security level for comparison: OpenSSL’s implementation of ECDH (on the `nistp256` curve) and RSA with a 3072-bit modulus. Results were measured on a single hardware hyper-thread on a 2.6GHz Intel Xeon E5 (Sandy Bridge); see [9] for details. Although some implementations included optimizations using the AVX2 instruction set, the computer used for measurements did not support AVX2.

In the table, Alice0 denotes Alice’s procedure for constructing her outgoing message, and Alice1 is her procedure for processing Bob’s incoming message and deriving the shared secret.

The NewHope protocol has the best computational performance of the post-quantum key exchange algorithms tested, even outperforming traditional RSA and ECDH. However, all structured lattice schemes (ring-LWE and NTRU) have larger communication than RSA and ECDH, around 2-8 KiB round-trip. Unstructured lattice schemes (LWE) also achieve good performance, on the order of 1 ms, but require even more communication, around 22 KiB round-trip. Supersingular isogeny Diffie–Hellman has much smaller keys (1 KiB round-trip, not much larger than RSA 3072), but orders of magnitude slower performance. (Note, however, that the AVX2 optimized implementation of SIDH was an order of magnitude faster than its C implementation). No code-based post-quantum protocol was included in the tests above. In particular, the implementation of Bernstein et al.’s “McBits” high-speed code-based cryptosystem [6] was not publicly available at the time of writing, but their paper reports speeds of 0.005ms (on a 3.4 GHz CPU) for decryption at the 128-bit quantum security level, but at the cost of 216 KiB public keys.

These trade-offs leave no clear post-quantum winner: the smallest key sizes come from SIDH but it has slow performance (though performance usually improves!); ring-LWE gives a decent tradeoff with fast performance and not-too-big keys; LWE’s performance remains good, and avoids the use of a structured lattice, but requires larger communication. Though these larger public keys may be too big for embedded devices, it should be remembered that the average webpage is over 1 MB: if we had to switch the Internet to post-quantum cryptography today, the communication costs from post-quantum key exchange would not be much more than an extra emoticon on a webpage.

3.5 From unauthenticated to authenticated key exchange

Both the BCNS15 and Frodo protocols are for unauthenticated key exchange: they assume the adversary is passive. Of course in practice one must achieve security against an active network adversary. Peikert [43] noted the challenges that are faced in securing LWE and ring-LWE based protocols against an active adversary, and Fluhrer [21] described an explicit attack on ring-LWE protocols that reuse ephemeral key shares against an active adversary. Peikert suggested the use of a transform such as the Fujisaki–Okamoto transform [23] which converts a passively secure (IND-CPA) key encapsulation mechanism (KEM) into an actively secure (IND-CCA) KEM. For integration with TLS, there is also the possibility of using signatures in a signed-DH-like protocol to first authenticate the keyshares; see [10].

4 Integrating post-quantum key exchange into TLS

All the post-quantum key exchange candidates explored in the previous section incur some penalty (either slower computation, or bigger communication, or both) compared to existing RSA or elliptic curve public key cryptography. It is therefore important to understand the impact of these penalties in a practical setting. Both the BCNS15 and Frodo papers integrate the corresponding key exchange scheme into the Transport Layer Security (TLS) protocol, the dominant protocol used securing Internet communications. In particular, they create new TLS version 1.2 ciphersuites which use traditional RSA or ECDSA certificates for signature, but use post-quantum key exchange to derive a shared secret, and then continue to use standard TLS authenticated encryption constructions (e.g., AES in GCM mode). (Due to the message flow in the TLS 1.2 handshake, the TLS server plays the role of “Alice” in the key exchange, and the TLS client plays the role of “Bob”.) This is achieved by modifying OpenSSL, a common open-source library for SSL/TLS, which is used by applications such as the Apache httpd web server for securing web server communication.

Hybrid ciphersuites. The experiments involving post-quantum ciphersuites in TLS also included hybrid ciphersuites, where the TLS handshake uses two key exchange algorithms: one post-quantum algorithm, and one traditional algorithm (in this case, ECDH). While the use of two key exchange algorithms does impact performance, it allows early adopters to retain the (current) security of traditional algorithms like ECDHE while obtaining (potential) security against quantum computers: since many post-quantum algorithms have had comparatively less cryptanalysis, there is an increased chance that parameter sizes for post-quantum algorithms will evolve more rapidly over the next few years in the face of new classical or quantum cryptanalytic advances. Interestingly, Google, in its recent NewHope experiment in Chrome, decided to use solely hybrid ciphersuites [12].

Security. As noted above, BCNS15 and Frodo were shown to be secure *unauthenticated* key exchange protocols, i.e., assuming a passive adversary. For security against an active adversary, we showed in the BCNS paper [10] how to achieve the standard security notion for TLS (“authenticated and confidential channel establishment” (ACCE) [27]) if the server signs both the client and server key share. Note that this would require reordering some of the messages in TLS. An alternative, as noted above, is to use a KEM transform to obtain an actively-secure key exchange protocol.

4.1 Performance of post-quantum key exchange in TLS

Table 2 (copied from [9]) shows the performance of a TLS-protected Apache web server using various key exchange mechanisms and signature schemes. It measures:

Ciphersuite		Connections/second			Connection time (ms)		Handshake
Key exchange	Signature	1B	10 KiB	100 KiB	w/o load	w/load	size (bytes)
ECDHE nistp256	ECDSA	1187	1088	961	14.2	22.2	1,264
	RSA	814	790	710	16.1	24.7	1,845
BCNS15	ECDSA	922	893	819	18.8	35.8	9,455
	RSA	722	716	638	20.5	36.9	9,964
NewHope	ECDSA	1616	1351	985	12.1	18.6	5,005
	RSA	983	949	771	13.1	20.0	5,514
NTRU EES743EP1	ECDSA	725	708	612	20.0	27.2	3,181
	RSA	553	548	512	19.9	29.6	3,691
Frodo Recomm.	ECDSA	923	878	843	18.3	31.5	23,725
	RSA	703	698	635	20.7	32.7	24,228
Hybrid ciphersuites							
BCNS15+ECDHE	ECDSA	736	728	664	23.1	37.7	9,595
	RSA	567	559	503	24.6	40.2	10,177
NewHope+ECDHE	ECDSA	1095	1017	776	16.5	25.2	5,151
	RSA	776	765	686	18.1	28.0	5,731
NTRU+ECDHE	ECDSA	590	578	539	22.5	34.3	3,328
	RSA	468	456	424	24.2	36.8	3,908
Frodo Rec.+ECDHE	ECDSA	735	701	667	22.9	36.4	23,859
	RSA	552	544	516	24.5	39.9	24,439

Table 2: **Performance of Apache httpd web server**, measured in connections per second, connection time in milliseconds, and handshake size in bytes. Table from [9]. All TLS ciphersuites used AES256-GCM authenticated encryption with SHA384 in the MAC and KDF. Note that different key exchange methods are at different security levels; see Table 1 for details.

- *throughput* (connections/second): number of connections per second at the server before server latency spikes, measured with requests of different payload sizes (1B, 10 KiB, 100 KiB);
- *handshake latency* (milliseconds): time from when client sends first TCP packet till client receives first application data packet, measured on an unloaded server and a loaded server (with sufficiently many connections to achieve 70% CPU load).

Performance was measured on a 4-CPU server with the same CPU as in Section 3.4. See [9] for the detailed methodology.

Unsurprisingly, the additional overhead of other cryptographic and network operations in a TLS connection mutes the performance differences between key exchange protocols. For example, while the standalone performance of NewHope is $9\times$ better than that of Frodo recommended, throughput of a NewHope-based ciphersuite is only $1.75\times$ better than Frodo recommended when the server returns 1 byte of application data, and the gap narrows further to just $1.12\times$ when the server returns 100 KiB of application data. Similarly, the latency of a Frodo-based ciphersuite is only $1.5\times$ slower than a NewHope-based ciphersuite. When hybrid ciphersuites are used, the performance difference between slow and fast post-quantum ciphersuites narrows even further.

Component	New files	Existing files	Lines of code*
Core ring-LWE library	6	0	896
Ring-LWE “wrapper” for OpenSSL	6	5	1229
SSL integration	0	12	914

Table 3: **Source code changes to add BCNS15 ring-LWE key exchange to OpenSSL.**

* Lines of code excludes Makefiles and automatically generated files, but includes comments and whitespace, and counts both lines added and deleted. Calculated from <https://github.com/dstebila/openssl-rlwekex/commit/f80719bf>.

5 Interlude: programming is hard

In the BCNS15 work on ring-LWE-based key exchange, we did a performance evaluation at two levels: the standalone cryptographic operations of the ring-LWE key exchange protocol, and its performance when run in the TLS protocol. The first is a fairly common practice in cryptographic research: implement your algorithms in C, then use some cycle counting or microsecond-accurate timing code to determine the runtime of your algorithms.

Evaluating performance in the TLS protocol is less common due in part to the difficulty of doing so: either one has to implement a network protocol from scratch (which is painful and usually not the main purpose of the research), or integrate the cryptographic algorithms into an existing cryptographic library, such as OpenSSL. These libraries are often quite complex. When we wanted to add our BCNS15 ring-LWE key exchange protocol to OpenSSL for testing purposes, we had to first “wrap” our core ring-LWE library inside of OpenSSL’s data structures inside the `crypto` directory, then modify OpenSSL’s `ssl` directory to make use of those new data structures. Table 3 shows the number of files and lines of code that were added or changed. While the core ring-LWE library consisted of only 6 (standalone) C files totalling just under 900 lines of code, integrating it into OpenSSL required touching 23 files and changing or adding another 2143 lines of code.

For the Frodo work on LWE-based key exchange, we again wanted to evaluate the performance of our algorithms both in a standalone setting and in the context of TLS, but we also wanted to compare with other post-quantum key exchange candidates. Writing 2100 lines of wrapper/integration code for each algorithm we wanted to add was an unappealing prospect. For the Frodo project, we developed an intermediate API that allowed us to more easily integrate different post-quantum key exchange algorithms into OpenSSL for performance comparison. This not-publicly-released intermediate API was the predecessor of and partial motivation for some of the features added to the Open Quantum Safe framework.

6 Open Quantum Safe: a software framework for post-quantum cryptography

The goal of our Open Quantum Safe (OQS) project (<https://openquantumsafe.org>) is to support the development and prototyping of quantum-resistant cryptography. OQS consists of two main lines of work: `liboqs`, an open source C library for quantum-resistant cryptographic algorithms; and prototype integrations into protocols and applications, including the widely used OpenSSL library.

As an example of where the OQS framework can assist with the grand challenge of moving

quantum-resistant cryptography towards reliable widespread deployment, consider a small- or medium-sized enterprise that understands the need to integrate quantum-resistant cryptography into its products. Perhaps their products protect information that requires long-term confidentiality. Perhaps their products will be deployed in the field for many years with no easy opportunity for changing the cryptographic algorithms later. Or perhaps they or their customers are worried about the small but non-negligible chance that today’s algorithms will be broken, by quantum computers or otherwise, much earlier than expected.

Whatever their reason for wishing to integrate quantum-safe cryptography into their products sooner rather than later, this would not be an easy path for them to take. In-house implementation of quantum-safe primitives requires advanced specialized expertise in order to understand the research literature, choose a suitable scheme, digest the new mathematics, choose suitable parameters, and develop robust software or hardware implementations. This is an enormous, expensive, and risky endeavour to undertake on one’s own, especially for a small- or medium-sized enterprise.

Commercially available alternatives, especially back in 2014 when this project started taking shape, were few, and also potentially problematic from a variety of perspectives: cost, patents, transparency, maintenance, degree of external scrutiny, etc.

Companies who would like to offer a quantum-safe option today do not have an easy or robust path for doing so.

OQS gives such organizations the option of prototyping an available quantum-resistant algorithm in their applications. Since these are still largely experimental algorithms that have not yet received the intense scrutiny of the global cryptographic community, our recommendation is to use one of the available post-quantum algorithms in a “hybrid” fashion with a standard algorithm that has received intense scrutiny with respect to classical cryptanalysis and robust implementation.

Since we fully expect that ongoing developments and improvements in the design, cryptanalysis, and implementation of quantum-safe algorithms, OQS is designed so improvements and changes in the post-quantum algorithm can be adopted without major changes to application software.

Organizations who do not wish or need to use open source in their products can still benefit from:

- reference implementations that will guide them in their own implementations
- benchmark information that will guide their choice of algorithm
- the ability to test alternatives in their products before deciding which algorithms to choose.

OQS was thus designed with the goal of both facilitating the prototyping and testing of quantum-resistant algorithms in a range of applications, and of driving forward the implementation, testing, and benchmarking of quantum-resistant primitives themselves.

The high-level architecture of the OQS software project is shown in Figure 3.

6.1 liboqs

liboqs (<https://github.com/open-quantum-safe/liboqs>) provides a common interface for key exchange and digital signature schemes, as well as implementations of a variety of post-quantum schemes. Some implementations are based on existing open source implementations, either adapting the implementation or putting a thin “wrapper” around the implementation. Other implementations have been written from scratch directly for the library. As of writing, liboqs includes key exchange based on:

- ring-LWE using the BCNS15 protocol (adaptation of existing implementation) [10];
- ring-LWE using the NewHope protocol (wrapper around existing implementation) [4];
- LWE using the Frodo protocol (adaptation of existing implementation) [9];

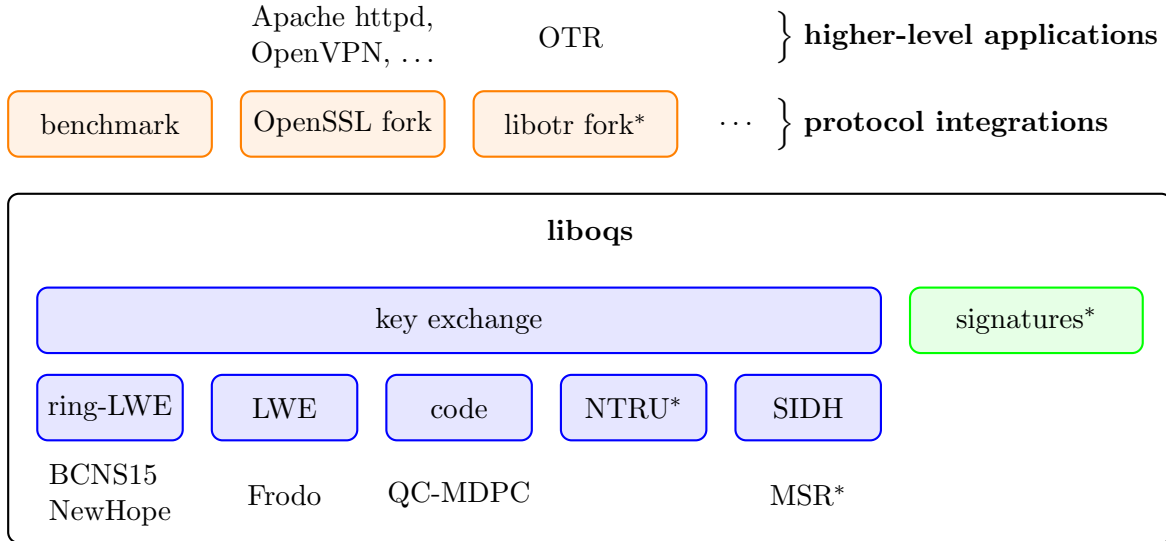


Figure 3: **Architecture of the Open Quantum Safe project.** * denotes future plans.

- error correcting codes – quasi-cyclic medium-density parity-check codes using the Niederreiter cryptosystem (new implementation).

liboqs also includes common routines available to all liboqs modules, including a common random number generator and various symmetric primitives such as AES and SHA-3.

liboqs includes a benchmarking program that enables runtime comparisons of all supported implementations. The library and benchmarking program build and have been tested on Mac OS X 10.11.6, macOS 10.12, and Ubuntu 16.04.1 (using clang or gcc), and Windows 10 (using Visual Studio).

6.2 Application/protocol integrations

The OQS project also includes prototype integrations into protocols and applications. Our first integration is into the OpenSSL library,³ which is an open source cryptographic library that provides both cryptographic functions (`libcrypto`) and an SSL/TLS implementation (`libssl`). OpenSSL is used by many network applications, including the popular Apache httpd web server and the OpenVPN virtual private networking software.

Our OpenSSL 1.0.2 fork (<https://github.com/open-quantum-safe/openssl>) integrates post-quantum key exchange algorithms from liboqs into OpenSSL’s `speed` command, and provides TLS 1.2 ciphersuites using post-quantum key exchange based on primitives from liboqs. For each post-quantum key exchange primitive supported by liboqs, there are ciphersuites with AES-128 or AES-256 encryption in GCM mode (with either SHA-256 or SHA-384, respectively), and authentication using either RSA or ECDSA certificates. (We use experimental ciphersuite numbers.)

Each of these four ciphersuites is also mirrored by another four *hybrid* ciphersuites which use both elliptic curve Diffie–Hellman (ECDHE) key exchange *and* the post-quantum key exchange primitive.

Our OpenSSL integration also includes *generic* ciphersuites. liboqs includes interfaces for each key exchange algorithm so it can be selected by the caller at runtime, but it also includes a generic

³<https://www.openssl.org>

interface that can be configured *at compile time*. Our OpenSSL integration does include ciphersuites for each individual key exchange algorithm in liboqs, but it also includes a set of ciphersuites that call the generic interface, which will then use whatever key exchange algorithm was specified at compile time. This means that a developer can add a new algorithm to liboqs and immediately prototype its use in SSL/TLS without changing a single line of code in OpenSSL, simply by using the generic OQS ciphersuites in OpenSSL and compiling liboqs to use the desired algorithm.

6.3 Case study: adding NewHope to liboqs and OpenSSL

As mentioned earlier, one of the goals of the Open Quantum Safe project is to make it easier to prototype post-quantum cryptography. It should be easy to add a new algorithm to liboqs, and then easy to use that algorithm in an application or protocol that already supports liboqs.

Recently, we added the NewHope ring-LWE-based key exchange to liboqs and our OpenSSL fork. It is interesting to compare the amount of work required to add NewHope to liboqs and our OpenSSL fork with the figures in Table 3 on adding BCNS15 directly to OpenSSL.

In liboqs, the wrapper around NewHope is 2 new files, totalling 163 lines of code, and requires 5 lines of code to be changed in 2 other files (plus changes in the Makefile).

As noted above, liboqs includes a “generic” key exchange method which can be hard-coded at compile time to any one of its implementations, and our OpenSSL fork already includes a “generic OQS” key exchange ciphersuite that calls liboqs’ generic key exchange method. Thus, once NewHope has been added to liboqs, it is possible to test NewHope in OpenSSL with zero changes to the OpenSSL fork via the generic key exchange method and recompiling. However, to explicitly add named NewHope ciphersuites to OpenSSL, we are able to reuse existing data structures, resulting in a diff that touches 10 files and totals 222 lines of code. Moreover, the additions can very easily follow the pattern from previous diffs,⁴ making adding a new OQS-based ciphersuite a 15-minute job.

7 Conclusion and outlook

The next few years will be an exciting time in the area of post-quantum cryptography. With the forthcoming NIST post-quantum project, and with continuing advances in quantum computing research, there will be increasing interest from government, industry, and standards bodies in understanding and using quantum-resistant cryptography. Lattice-based cryptography, in the form of the learning with errors and the ring-LWE problems, is particularly promising for quantum-resistant public key encryption and key exchange, offering high computation efficiency with reasonable key sizes. More cryptanalytic research will be essential to increase confidence in any standardized primitive. Since each post-quantum candidate to date has trade-offs between computational efficiency and communication sizes compared to existing primitives, it is also important to understand the how applications and network protocols behave when using different post-quantum algorithms. The Open Quantum Safe project can help rapidly compare post-quantum algorithms and prototype their use in existing protocols and applications, and experiments like Google’s use of NewHope in its Chrome Canary browser will give valuable information about how post-quantum cryptosystems behave in real-world deployments.

For cryptographers interested in designing new public key encryption, digital signature schemes, and key exchange protocols—for cryptanalysts looking to study new mathematical problems—

⁴<https://github.com/open-quantum-safe/openssl/commit/cb91c708> and <https://github.com/open-quantum-safe/openssl/commit/3a04b822>

for cryptographic engineers building new systems—and for standards bodies preparing for the future—exciting times lie ahead!

8 Acknowledgements

Research on LWE and ring-LWE based key exchange discussed in this paper includes joint work with Joppe W. Bos (NXP Semiconductors), Craig Costello (Microsoft Research), Léo Ducas (CWI), Ilya Mironov (Google), Michael Naehrig (Microsoft Research), Valeria Nikolaenko (Stanford University), and Ananth Raghunathan (Google) and was published as [9] and [10].

The Open Quantum Safe project grew out of discussions with a number of colleagues over the past few years, especially during early 2014, around the challenges of taking post-quantum cryptography closer to widespread practical application. These colleagues include: Scott Vanstone and Sherry Shannon Vanstone (Trustpoint); Matthew Campagna (Amazon Web Services); Alfred Menezes, Ian Goldberg, and Guang Gong (University of Waterloo); William Whyte and Zhenfei Zhang (Security Innovation); as well as the research colleagues in the paragraph above; we are grateful for all the valuable discussions. The Open Quantum Safe project has been supported in part by the Tutte Institute for Mathematics and Computing. Contributors to OQS are listed at <https://github.com/open-quantum-safe/liboqs/graphs/contributors> and as of writing include Jennifer Fernick, David Jao, Tancrede Lepoint, Shravan Mishra, Christian Paquin, Alex Parent, John Schanck, and Sebastian Verschoor.

D.S. thanks the chairs of SAC 2016 for the invitation to give the Stafford Tavares Invited Lecture at Canada’s annual cryptography conference.

References

- [1] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996.
- [2] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *29th ACM STOC*, pages 284–293. ACM Press, May 1997.
- [3] M. R. Albrecht, R. Player, and S. Scott. On the concrete hardness of learning with errors. Cryptology ePrint Archive, Report 2015/046, 2015. <http://eprint.iacr.org/2015/046>.
- [4] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange - a new hope. In *USENIX Security 2016*. USENIX Association, Aug. 2016. <https://eprint.iacr.org/2015/1092>.
- [5] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, Aug. 2009.
- [6] D. J. Bernstein, T. Chou, and P. Schwabe. McBits: Fast constant-time code-based cryptography. In G. Bertoni and J.-S. Coron, editors, *CHES 2013*, volume 8086 of *LNCS*, pages 250–272. Springer, Aug. 2013.
- [7] D. J. Bernstein, D. Hopwood, A. Hülsing, T. Lange, R. Niederhagen, L. Papachristodoulou, M. Schneider, P. Schwabe, and Z. Wilcox-O’Hearn. SPHINCS: Practical stateless hash-based signatures. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 368–397. Springer, Apr. 2015.
- [8] O. Blazy, C. Chevalier, L. Ducas, and J. Pan. Exact smooth projective hash function based on LWE. Cryptology ePrint Archive, Report 2013/821, 2013. <http://eprint.iacr.org/2013/821>.

- [9] J. W. Bos, C. Costello, L. Ducas, I. Mironov, M. Naehrig, V. Nikolaenko, A. Raghunathan, and D. Stebila. Frodo: Take off the ring! practical, quantum-secure key exchange from LWE. In *ACM CCS 2016*. ACM Press, Oct. 2016. <https://eprint.iacr.org/2016/659>.
- [10] J. W. Bos, C. Costello, M. Naehrig, and D. Stebila. Post-quantum key exchange for the TLS protocol from the ring learning with errors problem. In *2015 IEEE Symposium on Security and Privacy*, pages 553–570. IEEE Computer Society Press, May 2015.
- [11] J. W. Bos, M. Naehrig, and J. van de Pol. Sieving for shortest vectors in ideal lattices: a practical perspective. Cryptology ePrint Archive, Report 2014/880, 2014. <http://eprint.iacr.org/2014/880>.
- [12] M. Braithwaite. Google Security Blog: Experimenting with post-quantum cryptography, July 2016. <https://security.googleblog.com/2016/07/experimenting-with-post-quantum.html>.
- [13] J. Buchmann, E. Dahmen, and A. Hülsing. XMSS - a practical forward secure signature scheme based on minimal security assumptions. In B.-Y. Yang, editor, *PQCrypto 2011*, volume 7071 of *LNCS*, pages 117–129. Springer, 2011.
- [14] Y. Chen and P. Q. Nguyen. BKZ 2.0: Better lattice security estimates. In D. H. Lee and X. Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 1–20. Springer, Dec. 2011.
- [15] C. Costello, P. Longa, and M. Naehrig. Efficient algorithms for supersingular isogeny Diffie-Hellman. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 572–601. Springer, Aug. 2016.
- [16] R. Cramer, L. Ducas, and B. Wesolowski. Short Stickelberger class relations and application to Ideal-SVP. Cryptology ePrint Archive, Report 2016/885, 2016. <http://eprint.iacr.org/2016/885>.
- [17] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: an outlook. *Science*, 339(6124):1169–1174, 2013.
- [18] T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246 (Proposed Standard), Aug. 2008.
- [19] J. Ding, X. Xie, and X. Lin. A simple provably secure key exchange scheme based on the learning with errors problem. Cryptology ePrint Archive, Report 2012/688, 2012. <http://eprint.iacr.org/2012/688>.
- [20] L. D. Feo, D. Jao, and J. Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *J. Mathematical Cryptology*, 8(3):209–247, Sept. 2014.
- [21] S. Fluhrer. Cryptanalysis of ring-LWE based key exchange with key share reuse. Cryptology ePrint Archive, Report 2016/085, 2016. <http://eprint.iacr.org/2016/085>.
- [22] A. Fujioka, K. Suzuki, K. Xagawa, and K. Yoneyama. Practical and post-quantum authenticated key exchange from one-way secure key encapsulation mechanism. In K. Chen, Q. Xie, W. Qiu, N. Li, and W.-G. Tzeng, editors, *ASIACCS 13*, pages 83–94. ACM Press, May 2013.
- [23] E. Fujisaki and T. Okamoto. How to enhance the security of public-key encryption at minimum cost. In H. Imai and Y. Zheng, editors, *PKC'99*, volume 1560 of *LNCS*, pages 53–68. Springer, Mar. 1999.
- [24] B. Gong and Y. Zhao. Small field attack, and revisiting RLWE-based authenticated key exchange from Eurocrypt'15. Cryptology ePrint Archive, Report 2016/913, 2016. <http://eprint.iacr.org/2016/913>.
- [25] S. Gueron and F. Schlieker. Speeding up R-LWE post-quantum key exchange. Cryptology ePrint Archive, Report 2016/467, 2016. <http://eprint.iacr.org/2016/467>.
- [26] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring based public key cryptosystem. In *Algorithmic Number Theory (ANTS III)*, volume 1423 of *LNCS*, pages 267–288, 1997.
- [27] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DHE in the standard model. In R. Safavi-Naini and R. Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 273–293. Springer, Aug. 2012.

- [28] J. Katz and V. Vaikuntanathan. Smooth projective hashing and password-based authenticated key exchange from lattices. In M. Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 636–652. Springer, Dec. 2009.
- [29] T. Laarhoven. Sieving for shortest vectors in lattices using angular locality-sensitive hashing. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 3–22. Springer, Aug. 2015.
- [30] L. Lamport. Constructing digital signatures from a one-way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory, Oct. 1979.
- [31] R. Lindner and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In A. Kiayias, editor, *CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339. Springer, Feb. 2011.
- [32] P. Longa and M. Naehrig. Speeding up the number theoretic transform for faster ideal lattice-based cryptography. Cryptology ePrint Archive, Report 2016/504, 2016. <http://eprint.iacr.org/2016/504>.
- [33] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In H. Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, May 2010.
- [34] V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for ring-LWE cryptography. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, May 2013.
- [35] T. Matsumoto and H. Imai. Public quadratic polynomial-tuples for efficient signature-verification and message-encryption. In C. G. Günther, editor, *EUROCRYPT’88*, volume 330 of *LNCS*, pages 419–453. Springer, May 1988.
- [36] R. McEliece. A public-key cryptosystem based on algebraic coding theory. DSN Progress Report 42-44, January and February 1978.
- [37] R. C. Merkle. A certified digital signature. In G. Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 218–238. Springer, Aug. 1990.
- [38] M. Mosca. Cybersecurity in an era with quantum computers: Will we be ready? Cryptology ePrint Archive, Report 2015/1075, 2015. <http://eprint.iacr.org/2015/1075>.
- [39] National Security Agency. NSA Suite B cryptography: Cryptography today, Aug. 2015. https://www.nsa.gov/ia/programs/suiteb_cryptography/.
- [40] H. Niederreiter. Knapsack-type cryptosystems and algebraic coding theory. *Problems of Control and Information Theory. Problemy Upravleniya i Teorii Informacii*, 15:159–166, 1986.
- [41] H. Ong and C.-P. Schnorr. Signatures through approximate representation by quadratic forms. In D. Chaum, editor, *CRYPTO’83*, pages 117–131. Plenum Press, New York, USA, 1983.
- [42] C. Peikert. Some recent progress in lattice-based cryptography. In O. Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, page 72. Springer, Mar. 2009.
- [43] C. Peikert. Lattice cryptography for the Internet. In M. Mosca, editor, *PQCrypto 2014*, volume 8772 of *LNCS*, pages 197–219. Springer, 2014.
- [44] C. Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 10(4):283–424, 2016.
- [45] C. Peikert. How (not) to instantiate ring-LWE. Cryptology ePrint Archive, Report 2016/351, 2016. <http://eprint.iacr.org/2016/351>.
- [46] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In H. N. Gabow and R. Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
- [47] M. Schneider. Sieving for shortest vectors in ideal lattices. In A. Youssef, A. Nitaj, and A. E. Hassanien, editors, *AFRICACRYPT 13*, volume 7918 of *LNCS*, pages 375–391. Springer, June 2013.
- [48] P. W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *35th FOCS*, pages 124–134. IEEE Computer Society Press, Nov. 1994.

- [49] J. Zhang, Z. Zhang, J. Ding, M. Snook, and Ö. Dagdelen. Authenticated key exchange from ideal lattices. In E. Oswald and M. Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 719–751. Springer, Apr. 2015.