An Injectivity Analysis of CRYSTALS-Kyber and Implications on Quantum Security*

Xiaohui Ding[1], Muhammed F. Esgin[1,2], Amin Sakzad[1], and Ron Steinfeld[1]

[1] Faculty of Information Technology, Monash University, Australia
[2] CSIRO's Data61, Australia
xd31412718@gmail.com
{muhammed.esgin,amin.sakzad,ron.steinfeld}@monash.edu

Abstract. The One-Way to Hiding (O2H) Lemma is a central component of proofs of chosen-ciphertext attack (CCA) security of practical public-key encryption schemes using variants of the Fujisaki-Okamoto (FO) transform in the Quantum Random Oracle Model (QROM). Recently, Kuchta *et al.* (EUROCRYPT '20) introduced a new QROM proof technique, called *Measure-Rewind-Measure* (MRM), giving an improved variant of the O2H lemma, with a new security reduction that does not suffer from a square-root advantage security loss as in the earlier work of Bindel *et al.* (TCC '19).However, the FO transform QROM CCA security reduction based on the improved MRM O2H lemma still requires an injectivity assumption on the underlying CPA-secure determinstic public-key encryption scheme. In particular, the tightness of the concrete security reduction relies on a sufficiently small injectivity bound, and obtaining such bounds for concrete schemes was left as an open problem by Kuchta *et al.* (EUROCRYPT '20).
In this paper, we address the above problem by deriving concrete bounds on the injectivity of the deterministic CPA-secure variant of CRYSTALS-Kyber, the public-key encryption scheme selected for standardisation by the NIST Post-Quantum Cryptograpy (PQC) standardisation process. We evaluate our bounds numerically for the CRYSTALS-Kyber parameter sets, and show that the effect of injectivity on the tightness of the QROM CCA security of the Fujisaki-Okamoto transformed Kyber KEM is negligible, i.e. allows for a tight QROM CCA security reduction. Consequently, we give tightest QROM CCA security bounds to date for a simplified 'single hashing' variant of Kyber CCAKEM against attacks with low quantum circuit depth. Our bounds apply for all the Kyber parameter sets, based on the hardness of the Module Learning with Errors (MLWE) problem.

Keywords: post-quantum cryptography · CRYSTALS-Kyber · one-way to hiding · tight security

* A preliminary version of this paper has been presented in ACISP 2022 [8]. This extended and updated version of the paper contains improved and simplified injectivity bounds and several corrections.

1 Introduction

Post-quantum cryptography (PQC) has been considered crucial and constantly developed for the last two decades since the fast database search algorithm by Grover [11] and the fast integer factorization algorithm by Shor [18] on quantum computers were introduced. When a quantum processor with enough qubits is built, it will put many current public-key cryptosystems in danger. That is why in 2016, NIST announced the first round PQC standardization process [14]. Now the result of the third round competition has been announced on July 05, 2022 [23]. CRYSTALS-Kyber [7], which utilises module learning with errors (MLWE) as its underlying mathematical problem, has been selected to standardize for Public-Key Encryption (PKE)/Key Encapsulation Mechanism (KEM). There have been many applications of it in the real world. For example, it was integrated into the CIRCL cryptography library of Cloudflare [21] and is also supported as one of the post-quantum Transport Layer Security (TLS) protocols in Amazon Web Services (AWS) key management service [25]. Our paper will focus on Kyber and investigate the injectivity of its underlying deterministic PKE and its implications on the chosen-ciphertext attack security of the Kyber KEM.

CRYSTALS-Kyber uses cryptographic hash functions to achieve indistinguishable chosen-ciphertext attack (IND-CCA) security. We model classical (respectively, quantum) attacks on schemes using these hash functions in the Random Oracle Model (respectively, the Quantum Random Oracle Model, QROM). First defined by Bellare and Rogaway in 1993 [4], ROM gives the attacker a mechanism (oracle) $\mathcal{O}$ that takes input query $x \in \{0,1\}^*$ and generates random output $\mathcal{O}(x) \in \{0,1\}^n$. If query x has appeared before, then $\mathcal{O}$ will return the same result as the first output. This is in contrast to a QROM, $\mathcal{O}_q$, which was first introduced by Boneh *et al.* [6]. QROM replaces the query x and the output $\mathcal{O}$ with the qubit query $|\psi\rangle = \sum_i \alpha_i |\psi_i\rangle$ and the qubit output $\mathcal{O}_q |\psi\rangle = \sum_i \alpha_i |\mathcal{O}_q(\psi_i)\rangle$, where $\alpha_i \in \mathbb{C}$ are the complex coefficients of the superposition such that $\sum_i |\alpha_i|^2 = 1$.

A security reduction proof for a cryptographic scheme relates the security of the cryptographic scheme to the hardness of an underlying computational problem. It is desirable to have a tight security reduction, i.e. a reduction which guarantees that breaking the cryptographic scheme (with respect to some appropriate security notion) is almost as hard as solving the underlying computational problem, to ensure that there are no 'shortcut' attacks against the scheme that bypass solving the underlying hard computational problem. Assume that the success probability (or advantage) of an adversary $\mathcal{A}$ taking time t to break some security notion for a cryptographic scheme $\mathcal{C}$ is ϵ, and that a security reduction shows how to use $\mathcal{A}$ to construct an algorithm $\mathcal{B}$ for solving the underlying hard computational problem taking time t' and success probability (or advantage) ϵ'. Informally, if $t' \approx t'$ and $\epsilon' \approx \epsilon$, then the reduction is said to be tight. The concept is widely used in proving the security reductions from indistinguishable chosen-plaintext attack (IND-CPA) of a Public-Key Encryption (PKE) to indistinguishable chosen-ciphertext attack (IND-CCA) of a Key-Encapsulation

Mechanism (KEM). By applying Fujisaki-Okamoto (FO) transform [9, 10], the IND-CCA security of a KEM can be reduced from the IND-CPA security of its underlying PKE in the ROM. However, when the adversary has quantum access, all the proof techniques for classical ROM fail because the quantum query now can evaluate the random oracle in exponentially many points rather than a polynomial amount in classical ROM [6].

Providing tight security proof of FO transform under QROM has been the goal of several recent studies. In 2015, Unruh [24] introduced an approach called One Way to Hiding (O2H) lemma for (a tighter) security proof. Several variants of O2H have been developed to investigate different aspects in the advantage of an adversary with the help of various assumptions. Ambainis *et al.* [1] put forward a semi-classical O2H to mitigate the problem of measuring in QROM. Adapting this work, Bindel *et al.* [5] introduced another O2H variant called 'double-sided', which has some additional assumptions compared to the original O2H [24] but also has a tighter security proof in some particular parameters. Later, Kuchta *et al.* [13] introduced and applied a novel measure-rewind-measure technique for proving a double-sided O2H lemma to obtain a CCA security reduction for the $\mathsf{FO}^{\not\perp}$ transform [12] in the QROM that does not suffer a square-root loss of advantage. In [13], Corollary 4.7, they summarise the CCA advantage inequality as below:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{IND-CCA}}_{\mathsf{FO}^{\not\perp}(\mathsf{P},\mathsf{F},G,H)}(\mathcal{A}) \leq O(d_Q^2 \cdot \mathsf{Adv}^{\mathsf{IND-CPA}}_{\mathsf{P}}(\mathcal{B}_1) \\ + Q/|\mathcal{M}| + Qd_Q \cdot \delta + Q \cdot \sqrt{\eta} + \mathsf{Adv}^{\mathsf{PRF}}_{\mathsf{F}}(\mathcal{B}_2)),\end{aligned} \tag{1}$$

where $\mathcal{A}$ is the IND-CCA adversary against the KEM $\mathsf{FO}^{\not\perp}(\mathsf{P},\mathsf{F},G,H) = U^{\not\perp}(\mathsf{P}',\mathsf{F},H)$, obtained by applying the $\mathsf{U}^{\not\perp}$ Fujisaki-Okamoto (FO) transform, an implicit rejection variant of FO transform defined in [12] (see also [5]), to the randomised PKE scheme P, using hash functions G, H modelled as quantum random oracles, and a Pseudo-Random Function F. In this equation, $\mathcal{B}_1$ is the IND-CPA adversary against the randomised PKE scheme P, and $\mathcal{B}_2$ is the Pseudo-Random Function (PRF) adversary against F. Also, in the above equation, Q denotes the total number of attacker QROM queries, d_Q denotes the attacker QROM query depth, δ is the decryption failure error, $|\mathcal{M}|$ is the size of the message space, and η is the injectivity parameter (non-injectivity probability) of the underlying deterministic PKE scheme $\mathsf{P}' := T(\mathsf{P}, G)$ scheme obtained from P by deriving the encryption randomness by applying the hash function G to the input message. It was left as an open problem in [13] to compute a bound on the injectivity parameter η of concrete schemes in the NIST PQC process, namely a bound on the collision probability of the existence of two different messages generating the same ciphertext. To achieve the goal of a tight reduction at a λ-bit security level guarantee for small d_Q, we would like the bound on the adversary advantage on the right-hand side of the above equation to be $O(Q \cdot 2^{-\lambda})$ for small d_Q. Due to the square-root term $\sqrt{\eta}$, this implies that we need injectivity bound to satisfy $\eta = O(2^{-2\lambda})$, (or equivalently, $\sqrt{\eta} = O(2^{-\lambda})$).

1.1 Our Contribution

Following the work done by Kuchta *et al.* [13], we initiate the investigation of the injectivity of concrete schemes. In particular, we present concrete analytical and numerical upper bounds on η-injectivity of the deterministic variant of CRYSTALS-Kyber, the public-key encryption scheme selected for standardisation by the NIST Post-Quantum Cryptograpy (PQC) standardisation process. Our results give the tightest concrete QROM security guarantees to date for the QROM CCA security of a variant of CRYSTALS-Kyber, based on MLWE hardness assumption. Our contributions are summarized as follows:

- **Concrete Analytical Injectivity Bound:** We divide the injectivity analysis into two parts contributed by centred binomial distribution and module short integer solution (MSIS) problem. We then give a concrete analytical bound (Theorem 2) on the injectivity of CRYSTALS-Kyber PKE/KEM as a combination of the separate parts.
- **Numerical Injectivity Bounds:** Using our analytical results, we compute numerical injectivity bounds for the Kyber parameter sets, and in particular, the square root of the injectivity term $\sqrt{\eta}$ in (1). We obtain, for Kyber512, $\sqrt{\eta} = 2^{-177}$, for Kyber768, $\sqrt{\eta} = 2^{-467}$, and for Kyber1024, $\sqrt{\eta} = 2^{-844}$. These numerical results, for all three parameter sets, satisfy our requirement $\sqrt{\eta} \leq 2^{-\lambda}$, which implies that the effect of injectivity on the tightness of the QROM CCA security of the Fujisaki-Okamoto transformed Kyber KEM is negligible, i.e. allows for a tight QROM CCA security reduction.
- **Numerical QROM CCA Security Bounds:** We analyse the concrete implications of our numerical injectivity bounds on the INDCCA security of the Kyber CCAKEM, using the results of [13], and obtain the tightest QROM CCA security bounds to date for Kyber CCAKEM against attacks with low quantum circuit depth, for the Kyber parameter sets. Our concrete bounds are stated for a simplified 'single hashing' model of Kyber CCAKEM (see Sec. 4.2 for further discussion).

2 Preliminaries

Rings, Matrices and Vectors. We use R to represent $\mathbb{Z}[X]/(X^n+1)$ and R_q to represent $\mathbb{Z}_q[X]/(X^n+1)$. The degree n of the monic polynomial is fixed to 256 in Kyber. Matrices and vectors are represented as bold upper-case and lower-case letters, respectively. We use $\mathbf{v}^T$ to represent the transpose of $\mathbf{v}$. We also set $[\beta] = \{-\beta, -\beta+1, \ldots, 0, \ldots, \beta-1, \beta\} \subseteq \mathbb{Z}_q$, with $\beta \geq 0$, to represent a symmetrical integer set.

Norm and Cardinality. For an element $w \in \mathbb{Z}_q$, we set l_∞-norm of w to be $\|w\|_\infty = \left|w \bmod^{\pm} q\right|$, where $\bmod^{\pm} q$ is the modulo operation that takes w to the range $\left(-\frac{q}{2}, \frac{q}{2}\right]$ (resp. $\left[-\frac{q-1}{2}, \frac{q-1}{2}\right]$) for an even(resp. odd) q. For a polynomial element $w = w_0 + w_1 X + \cdots + w_{n-1}X^{n-1} \in R$, the l_∞ and l_2 norm can be defined as the followings:

$$\|w\|_\infty = \max_i \|w_i\|_\infty, \quad \|w\| = \sqrt{\|w_0\|_\infty^2 + \cdots + \|w_{n-1}\|_\infty^2}.$$

For vector $\mathbf{w} = (w_1, \ldots, w_k) \in R^k$, the norms are defined as:

$$\|\mathbf{w}\|_\infty = \max_{1 \le i \le k} \|w_i\|_\infty, \|\mathbf{w}\| = \sqrt{\|w_1\|^2 + \cdots + \|w_k\|^2}.$$

As for a finite set $S \subseteq R^k$, we define $|S|$ as the cardinality of S, and we also have:

$$\|S\|_\infty = \max_{\mathbf{w} \in S} \|\mathbf{w}\|_\infty, \quad \|S\| = \max_{\mathbf{w} \in S} \|\mathbf{w}\|.$$

Sampling. Let $\mathcal{X}$ be a set or a probability distribution. Then $x \leftarrow \mathcal{X}$ represents that value x is uniformly sampled from this set or is sampled from this distribution. For a polynomial $f(x) \in R_q$ or a vector of such polynomials, this notation is defined coefficient-wise. Particularly, we denote β_η as the central binomial distribution over $\mathbb{Z}$ with parameter η (see Def. 2) and we extend it to a distribution over R_q by sampling each polynomial integer coefficient independently.

Rounding. Let $x \in \mathbb{R}$ be a real number, then $\lceil x \rfloor$ means rounding to the closet integer with ties rounded up. We also use $\lceil x \rceil$ to represent round-up and $\lfloor x \rfloor$ as rounding down.

Compress and Decompress. Let $x \in \mathbb{Z}_q$ and $d \in \mathbb{Z}$ be such that $d < \lceil \log_2(q) \rceil$. Adapted from [7], the Compress and Decompress functions are:

$$\begin{aligned} \mathsf{Compress}_q(x, d) &= \left\lceil \left(2^d/q\right) \cdot x \right\rfloor \bmod^+ 2^d, \\ \mathsf{Decompress}_q(x, d) &= \left\lceil \left(q/2^d\right) \cdot x \right\rfloor, \end{aligned}$$

where $\bmod^+$ maps an element to the range $[0, 2^d)$.

2.1 Injectivity

Adapted from Definition 6 of [5] and Definition 4.3 of [13], the injectivity that we will investigate is defined as below:

Definition 1 (Injectivity of a dPKE [5, 13]). *Let $\eta \geqslant 0$. A dPKE scheme $\mathsf{P} = (\mathsf{KeyGen}, \mathsf{Encr}, \mathsf{Decr})$ using a random oracle G is η-injective if*

$$\Pr\left(\mathsf{Encr}(\mathsf{pk}, \cdot) \textit{ is not injective: } (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda), G \xleftarrow{\$} \mathcal{G}\right) \le \eta,$$

where $G \xleftarrow{\$} \mathcal{G}$ is sampling a random element G uniformly from a finite set $\mathcal{G}$ of random functions, and a dPKE means PKE that has a deterministic encryption scheme.

2.2 CRYSTALS-Kyber Scheme

The PQC scheme that we are going to investigate is CRYSTALS-Kyber [7]. Since injectivity is defined for a deterministic PKE (see Def. 1), the variant of Kyber we describe here and study is the deterministic version of Kyber obtained by applying the T part of the Fujisaki-Okamoto CCA transform to the randomized Kyber encryption algorithm, i.e. the encryption algorithm randomness r is derived deterministically from the message m by hashing with a function G which we model in our injectivity analysis as a random oracle (see Sec. 4.2 for further details on the Fujisaki-Okamoto Transform). Furthermore, we use the CBD distribution (see Def. 2) parameters η_1, η_2 as defined in the latest version of the CRYSTALS-Kyber NIST PQC Round 3 specifications document at the time of writing [3]. Let q, k, n, d_t, d_u, d_v be positive integer parameters and $\mathcal{M}$ denotes the message space with $n = 256$ bit messages. We put the PKE algorithms of Kyber here as a reference.

Sam**.** Let x be a bit string and S be a distribution taking x as the input, then $y \sim S := \mathsf{Sam}(x)$ represents that the output y generated by distribution S and input x can be extended to any desired length. We remark that in our injectivity analysis we model Sam as a random oracle with the indicated output distributions.

Algorithm 1 $\mathsf{Kyber.CPA.KeyGen}(1^\lambda)$: key generation [7], pg.5, Algorithm 1

$\rho, \sigma \leftarrow \{0,1\}^{256}$
$\mathbf{A} \sim R_q^{k \times k} := \mathsf{Sam}(\rho)$
$(\mathbf{s}, \mathbf{e}) \sim \beta_{\eta_1}^k \times \beta_{\eta_1}^k := \mathsf{Sam}(\sigma)$
$\mathbf{t} := \mathsf{Compress}_q(\mathbf{As} + \mathbf{e}, d_t)$
return $(pk := (\mathbf{t}, \rho), sk := \mathbf{s})$

Algorithm 2 $\mathsf{Kyber.CPA.Enc}$ $(pk = (\mathbf{t}, \rho), m \in \mathcal{M})$ [7], pg. 5, Algorithm 2

$r := G(m) \leftarrow \{0,1\}^{256}$
$\mathbf{t} := \mathsf{Decompress}_q(\mathbf{t}, d_t)$
$\mathbf{A} \sim R_q^{k \times k} := \mathsf{Sam}(\rho)$
$(\mathbf{r}, \mathbf{e_1}, e_2) \sim \beta_{\eta_1}^k \times \beta_{\eta_2}^k \times \beta_{\eta_2}^k := \mathsf{Sam}(r)$
$\mathbf{u} := \mathsf{Compress}_q(\mathbf{A}^T\mathbf{r} + \mathbf{e}_1, d_u)$
$v := \mathsf{Compress}_q(\mathbf{t}^T\mathbf{r} + e_2 + \lceil \frac{q}{2} \rfloor \cdot m, d_v)$
return $c := (\mathbf{u}, v)$

Algorithm 3 $\mathsf{Kyber.CPA.Dec}(sk = \mathbf{s}, c = (\mathbf{u}, v))$ [7], pg.5, Algorithm 3

$\mathbf{u} := \mathsf{Decompress}_q(\mathbf{u}, d_u)$
$v := \mathsf{Decompress}_q(v, d_v)$
return $\mathsf{Compress}_q(v - \mathbf{s}^T\mathbf{u}, 1)$

2.3 Methodologies and Techniques

2.3.1 Operations on Probability Adapted from [15], we present some techniques for calculating probability as the multiplication of polynomials.

Law Convolution. Suppose A and B are random variables over $[\alpha]$ and $[\beta]$, respectively. Let a_i, b_j be the probability of A, B being equal to i, j for all $i \in [\alpha]$ and $j \in [\beta]$. Then we generate two polynomials

$$A(X) = \sum_{i=-\alpha}^{\alpha} a_i X^i, B(X) = \sum_{j=-\beta}^{\beta} b_j X^j,$$

to represent the probability of all possible outcomes of A and B. Now, define

$$C(X) = A(X) \cdot B(X) = \sum_{k=-(\alpha+\beta)}^{\alpha+\beta} c_k X^k,$$

to be the product of $A(X)$ and $B(X)$, where $k = i + j$ for each i, j. One can observe that the coefficient c_k is the probability of the sum of two independent random variables A and B is equal to k, i.e.,

$$\Pr(A + B = k) = C_k.$$

Thus, $C(X)$ in fact, can be used to represent the probability distribution of $A + B$. If we want to investigate the probability of independent multivariate, we can simply repeat the multiplication.

Law Product. Now, we want to calculate the probability of the product of two independent random variables, A and B. Let $A(X)$ and $B(X)$ be the polynomials as above. Then, define

$$D(X) = \sum_{i=-\alpha}^{\alpha} \sum_{j=-\beta}^{\beta} a_i b_j X^{ij} = \sum_{k=-\alpha\beta}^{\alpha\beta} d_k X^k$$

to be the law product of the two distributions, representing the probability distribution of $A \cdot B$.

Union Bound. Let $\{A_1, A_2, \ldots\}$ be a finite set of events (which are not necessarily independent), then the probability of at least one happening is less or equal to the summarized probability of every event as described below:

$$\Pr\left(\bigcup_{i=1}^{\infty} A_i\right) \leq \sum_{i=1}^{\infty} \Pr(A_i).$$

This is also called Boole's inequality.

2.3.2 Centred Binomial Distribution The coefficient of x^k in the binomial expansion $(x+1)^{\eta_0}$ is given by $\binom{\eta_0}{k} = \frac{\eta_0!}{k!(\eta_0-k)!}$. When we investigate the case

$$\left(\sqrt{x} + \frac{1}{\sqrt{x}}\right)^{2\eta_0} = \sum_{i=0}^{2\eta_0} \binom{2\eta_0}{i} x^{(-\eta_0+i)},$$

we can let $k = -\eta_0 + i$ and switch the right hand side to $\sum_{k=-\eta_0}^{\eta_0} \binom{2\eta_0}{\eta_0+k} x^k$. Therefore, a centred binomial distribution is defined as below:

Definition 2 (Centred Binomial Distribution). *A discrete random variable X is said to have centred binomial distribution β_{η_0} with parameter η_0, if it follows probability mass function below:*

$$f(k;\eta_0) = \Pr(X = k) = \frac{(2\eta_0)!}{(\eta_0+k)!(\eta_0-k)!} \cdot 2^{-2\eta_0}.$$

If we represent the probability of all outcomes of X in one polynomial:

$$\mathrm{P}(X) = \sum_{k=-\eta_0}^{\eta_0} \frac{(2\eta_0)!}{(\eta_0+k)!(\eta_0-k)!} \cdot 2^{-2\eta_0} \cdot X^k,$$

the polynomial then can be used to calculate the probability of the difference between 2 independent random variables X_1, X_2 being some value by combining techniques in Section 2.3.1. Let $X_1, X_2 \sim \beta_{\eta_0}$, the polynomial that represents the probability distribution of $X_1 - X_2$ can be written as:

$$\mathrm{P}(X_1 - X_2) = \mathrm{P}(X_1) \cdot \mathrm{P}(X_2) = \sum_{k=-2\eta_0}^{2\eta_0} p_k X^k. \tag{2}$$

2.3.3 Module Short Integer Solution Problem The short integer solution (SIS) problem can be briefly described as finding a short vector in a random lattice. When the lattice is defined on a module polynomial ring R_q, we call finding a short vector in such a lattice a module short integer solution (MSIS) problem. We adopt the following result (Theorem 1.1 and Corollary 3.9 of [17], which extends and tightens the previous bounds of [19,20]), giving a probabilistic lower bound on the norm of the shortest non-zero vector of such lattices.

Theorem 1 (Adapted from [17], Theorem 1.1 and Corollary 3.9). *Denote $S_\alpha := \{y \in R_q : \|y\|_\infty \leq \alpha\}$ and let $l, k, \alpha_1, \alpha_2 \in \mathbb{N}$, q prime, and assume that the quotient polynomial of R splits as follows mod q, that is $X^n + 1 = \prod_{i \leq d} f_i(X) \bmod q$, where $f_1(X), \ldots, f_d(X)$ denote irreducible factors in $\mathbb{Z}_q[X]$, each of degree n/d. Also, for $i = 1, \ldots, d$, define $W_i \subseteq R_q$ to be a set of polyno-*

mials such that $\forall u, v \in W_i$, $|\mathsf{Zero}(u-v)| < i$. *Then*

$$\Pr_{\boldsymbol{A} \leftarrow R_q^{k \times l}} \left[\exists (\boldsymbol{z_1}, \boldsymbol{z_2}) \in S_{\alpha_1}^l \setminus \{\boldsymbol{0}\} \times S_{\alpha_2}^k : \boldsymbol{A z_1} + \boldsymbol{z_2} = \boldsymbol{0}\right] \leq \frac{|S_{\alpha_1}|^l \cdot |S_{\alpha_2}|^k}{q^{nk}} + \sum_{i=1}^{e} \frac{\binom{d}{i} \cdot \left|S_{\alpha_1 + \|W_i\|_\infty}\right|^l \cdot \left|S_{\alpha_2 + \|W_i\|_\infty}\right|^k}{|W_i|^{l+k} \cdot q^{nk(1-i/d)}},$$

where e *is the largest integer such that* $\alpha_1 \sqrt{n} \geq q^{e/d}$, *and for* $y \in R_q$, $|\mathsf{Zero}(y)|$ *is the size of the set* $\mathsf{Zero}(y)$ *defined as follows:*

$$\mathsf{Zero}\left(y\right) := \{i : y \equiv 0 \left(\mathrm{mod}\left(f_i\left(X\right), q\right)\right)\}.$$

3 Theoretical Bounds for CRYSTALS-Kyber

Now we give details of calculating the η-injectivity of Kyber. We remind the reader that our injectivity result applies to the deterministic variant of Kyber described in Sec. 2.2, which is obtained from the randomized Kyber scheme by using the T transform with a random oracle G, i.e. deriving the encryption scheme randomness by applying G to the input message (this is the deterministic version of Kyber used in the Fujisaki-Okamoto CCAKEM variant of Kyber; see Sec. 4.2 for more details). Our injectivity analysis also models the Kyber probability distribution sampling algorithms Sam as random oracles that output the ideal distributions as indicated in Algorithms 1 and 2 in Sec. 2.2.

3.1 Main Result

We first give a theorem for Kyber injectivity as below and later demonstrate some essential lemmas for calculating the final equation.

Theorem 2 (η-injectivity of Kyber). *Let positive integers* $q, k, n, d, \eta_1, \eta_2$, d_u, d_v *represent the parameters of Kyber (see Sec. 2.2). Namely* q, k, n *represent the prime modulus, module rank, and ring* R *dimension, respectively, and* d *is the number of splitting factors of* R_q *(see Theorem 1). The parameters* η_1, η_2 *are the sampling parameters of centered binomial distribution and* $1 \leq d_u, d_v < \lceil \log_2(q) \rceil$ *are the scaling parameters of compress function. The injectivity (Definition 1) of Kyber is upper bounded by:*

$$\begin{aligned} \eta \leq & \binom{2^n}{2} \cdot \left(r_0 \cdot \sum_{j \in [\gamma_u]} e_j\right)^{nk} \\ & + \left\{ \left(\frac{(4\eta_1 + 1)(4\eta_2 + 2\gamma_u + 1)}{q} \right)^{nk} + \sum_{i=1}^{e} \binom{d}{i} \cdot \frac{1}{|W_i|^{2k}} \right. \\ & \left. \cdot \left\{ \frac{(4\eta_1 + 2\|W_i\|_\infty + 1)(4\eta_2 + 2\gamma_u + 2\|W_i\|_\infty + 1)}{q^{(1-i/d)}} \right\}^{nk} \right\}, \end{aligned} \tag{3}$$

where $e = \left\lfloor d \cdot \log_q(2\eta_1\sqrt{n}) \right\rfloor$, r_0 *and* e_j *are from* (7), (8), *respectively. The set* W_i *is constructed from* (10). *The integers* γ_u, γ_v *are defined as* $\gamma_u := \left\lfloor \frac{q}{2^{d_u}} \right\rfloor$ *and* $\gamma_v := \left\lfloor \frac{q}{2^{d_v}} \right\rfloor$.

Proof. Let $m, m' \in \mathcal{B}^{32}$ denote a pair of distinct 32-byte (256-bit) messages. According to Algorithm 2 and Definition 1, we want to find the upper bound of the probability that at least one such pair distinct messages m, m'exists such that the corresponding Kyber ciphertexts $c(m), c(m')$ for them are the same. Thus, let

$$\eta = \Pr\left(\cup_{m \neq m'} c\left(m\right) = c\left(m'\right)\right).$$

For each distinct message pair m, m', we denote the corresponding ciphertexts by $c(m) = (\mathbf{u}, v)$ and $c(m') = (\mathbf{u}', v')$, respectively. We first observe that in the above non-injectivity union event $\cup_{m \neq m'} c\left(m\right) = c\left(m'\right)$ (and all the subsequent union events analyzed below), it is sufficient to consider distinct message pairs $m \neq m'$ with $G(m) \neq G(m')$. This is because if $G(m) = G(m')$, then non-injectivity is impossible. Indeed, if $G(m) = G(m')$ then $(\mathbf{r}, e_2) = (\mathbf{r}', e_2')$, $v = \mathsf{Compress}_q(x, d_v)$, $v' = \mathsf{Compress}_q(x', d_v)$ where $x := \mathbf{t}^T\mathbf{r} + e_2 + \left\lceil \frac{q}{2} \right\rfloor \cdot m$ and $x' := \mathbf{t}^T\mathbf{r}' + e_2' + \left\lceil \frac{q}{2} \right\rfloor \cdot m'$. This would therefore imply that $x - x' = \left\lceil \frac{q}{2} \right\rfloor \cdot (m - m')$ has a coordinate with absolute value $\geq \left\lceil \frac{q}{2} \right\rfloor$, which implies using Lemma 1 below and $d_v \geq 1$ that $v \neq v'$ and hence non-injectivity is impossible. Therefore, from here onwards, in all the events over m, m' analyzed below, the event is implicitly over message pairs $m \neq m'$ such that $G(m) \neq G(m')$, but we do not write the restriction $G(m) \neq G(m')$ for the brevity of notation.

Next, we simplify the non-injectivity probability analysis by analyzing non-injectivity for the $\mathbf{u}$ ciphertext part only (this turns out to be sufficient to obtain good bounds in practice, see our computed numerical bounds in Sec. 4). Namely, the bound on η can be written as:

$$\begin{aligned} \eta &= \Pr\left(\cup_{m \neq m'} c(m) = c(m')\right) \\ &\leq \Pr\left(\cup_{m \neq m'} \Delta\mathbf{u} = \mathbf{u} - \mathbf{u}' = \mathbf{0}\right), \end{aligned}$$

where

$$\mathbf{u} - \mathbf{u}' := \mathsf{Compress}_q\left(\mathbf{A}^T\mathbf{r} + \mathbf{e}_1, d_u\right) - \mathsf{Compress}_q\left(\mathbf{A}^T\mathbf{r}' + \mathbf{e}_1', d_u\right). \quad (4)$$

We now investigate the injectivity brought up by the compress function. By applying Lemma 1, we get the coefficient-wise equation below:

$$\Delta\mathbf{u} = 0 \Longrightarrow \left\|\mathbf{A}^T\mathbf{r} + \mathbf{e}_1 - \mathbf{A}^T\mathbf{r}' + \mathbf{e}_1'\right\|_\infty = \left\|\mathbf{A}^T\Delta\mathbf{r} + \Delta\mathbf{e}_1\right\|_\infty < \frac{q}{2^{d_u}},$$

where $\Delta\mathbf{r} = \mathbf{r} - \mathbf{r}'$ and $\Delta\mathbf{e}_1 = \mathbf{e}_1 - \mathbf{e}_1'$. It is also noticed that the coefficients are all integers. Thus, we get:

$$\left\|\mathbf{A}^T\Delta\mathbf{r} + \Delta\mathbf{e}_1\right\|_\infty < \frac{q}{2^{d_u}} \Longrightarrow \mathbf{A}^T\Delta\mathbf{r} + \Delta\mathbf{e}_1 \in \left[\gamma_u\right],$$

where $\gamma_u := \left\lfloor \frac{q}{2^{d_u}} \right\rfloor$, and the equation is also coefficient-wise. We can further split these relationships into two situations by letting $\Delta\mathbf{r} = \mathbf{0}$ or $\Delta\mathbf{r} \neq \mathbf{0}$, which defines two parts of probability:

$$P_{CBD} := \max_{m \neq m'} \Pr\left(\mathbf{A}^T \Delta\mathbf{r} + \Delta\mathbf{e}_1 \in [\gamma_u] \cap \Delta\mathbf{r} = \mathbf{0}\right)$$

and

$$P_{MSIS} := \Pr\left(\cup_{m \neq m'} \mathbf{A}^T \Delta\mathbf{r} + \Delta\mathbf{e}_1 \in [\gamma_u] \cap \Delta\mathbf{r} \neq \mathbf{0}\right).$$

By adapting Lemma 2, we have:

$$P_{CBD} = \left(r_0 \cdot \sum_{j \in [\gamma_u]} e_j\right)^{nk},$$

where r_0 is the probability of single coefficient in $\Delta\mathbf{r}$ being 0, and e_j is the probability of single coefficient in $\Delta\mathbf{e}_1$ having the value j, for all j in range $[\gamma_u]$.

And applying Lemma 3 gives:

$$\begin{aligned} P_{MSIS} \leq & \left(\frac{(4\eta_1 + 1)(4\eta_2 + 2\gamma_u + 1)}{q}\right)^{nk} + \sum_{i=1}^{e} \binom{d}{i} \cdot \frac{1}{|W_i|^{2k}} \\ & \cdot \left\{\frac{(4\eta_1 + 2\|W_i\|_\infty + 1) \cdot (4\eta_2 + 2\gamma_u + 2\|W_i\|_\infty + 1)}{q^{(1-i/d)}}\right\}^{nk}. \end{aligned}$$

Finally, we combine CBD part and MSIS part to derive the theoretical bound of Kyber injectivity:

$$\begin{aligned} \eta &= \Pr\left(\cup_{m \neq m'} c(m) = c(m')\right) \\ &\leq \Pr\left(\cup_{m \neq m'} \Delta\mathbf{u} = \mathbf{u} - \mathbf{u}' = \mathbf{0}\right) \\ &= \Pr\left(\cup_{m \neq m'} \mathbf{A}^T \Delta\mathbf{r} + \Delta\mathbf{e}_1 \in [\gamma_u] \cap \Delta\mathbf{r} = \mathbf{0}\right) \\ &\qquad + \Pr\left(\cup_{m \neq m'} \mathbf{A}^T \Delta\mathbf{r} + \Delta\mathbf{e}_1 \in [\gamma_u] \cap \Delta\mathbf{r} \neq \mathbf{0}\right) \\ &\leq M \cdot \max_{m \neq m'} \Pr\left(\mathbf{A}^T \Delta\mathbf{r} + \Delta\mathbf{e}_1 \in [\gamma_u] \cap \Delta\mathbf{r} = \mathbf{0}\right) \\ &\qquad + \Pr\left(\cup_{m \neq m'} \mathbf{A}^T \Delta\mathbf{r} + \Delta\mathbf{e}_1 \in [\gamma_u] \cap \Delta\mathbf{r} \neq \mathbf{0}\right) \\ &= M \cdot P_{CBD} + P_{MSIS}, \end{aligned}$$

where $M = \binom{2^n}{2}$, which results in (3). □

3.2 Associated Lemmas and their Proofs

Now we demonstrate several lemmas to help us compute each component of η-injectivity.

Lemma 1 ($\mathsf{Compress}_q$ Equality Condition). *Let $x, x' \in \mathbb{Z}_q$ and $d \in \mathbb{Z}$ be such that $d < \lceil \log_2(q) \rceil$, then we have:*

$$\mathsf{Compress}_q(x, d) - \mathsf{Compress}_q(x', d) = 0 \Rightarrow |x - x'| < \frac{q}{2^d}. \tag{5}$$

Proof. We first look at the rounding of x, x'. If $\lceil x \rfloor - \lceil x' \rfloor = 0$, then we have $|x - x'| < 1$. If x, x' are multiplied by a constant c, this can be further written as:

$$\lceil cx \rfloor - \lceil cx' \rfloor = 0 \Rightarrow |cx - cx'| < 1 \Longleftrightarrow |x - x'| < \frac{1}{|c|}.$$

Now, we look at the equation defined by [3]:

$$\mathsf{Compress}_q(x, d) = \left\lceil \left(2^d/q\right) \cdot x \right\rfloor \bmod^+ 2^d.$$

Since $\left\lceil \left(2^d/q\right) \cdot x \right\rfloor$ maps $x \in \mathbb{Z}_q$ to $\mathbb{Z}_{2^d}$, the $\bmod^+ 2^{\mathrm{d}}$ operation only shifts the representatives of $\mathbb{Z}_{2^d}$ to $\{0, \ldots, 2^d - 1\}$, which means it does not affect the distance between two scaled elements. Therefore, we can safely remove $\bmod^+ 2^{\mathrm{d}}$ when calculating the difference between two compressed inputs:

$$\begin{aligned}
\Delta\mathsf{Compress}_q(x, d) &:= \mathsf{Compress}_q(x, d) - \mathsf{Compress}_q(x', d) \\
&= \left\lceil \left(2^d/q\right) \cdot x \right\rfloor \bmod^+ 2^d - \left\lceil \left(2^d/q\right) \cdot x' \right\rfloor \bmod^+ 2^d \\
&= \left\lceil \left(2^d/q\right) \cdot x \right\rfloor - \left\lceil \left(2^d/q\right) \cdot x' \right\rfloor.
\end{aligned}$$

This is the condition for two compressed inputs being equal as in (5). These results can be generalised to vectors in a component-wise fashion. □

Lemma 2 (CBD Injectivity). *Let $\mathbf{r}, \mathbf{r}' \in R_q^k$ be samples from β_{η_1} and $\mathbf{e}_1, \mathbf{e}_1' \in R_q^k$ be samples from β_{η_2}. Let d_u be a positive integer that $d_u < \lceil \log_2(q) \rceil$. The probability P_{CBD} is given by:*

$$P_{CBD} = \left(r_0 \cdot \sum_{j \in [\gamma_u]} e_j \right)^{nk}, \tag{6}$$

where $\gamma_u := \left\lfloor \frac{q}{2^{d_u}} \right\rfloor$, and

$$\sum_{j=-2\eta_1}^{2\eta_1} r_j X^j = \left(\sum_{i=-\eta_1}^{\eta_1} \frac{(2\eta_1)!}{(\eta_1 + i)!(\eta_1 - i)!} \cdot 2^{-2\eta_1} \cdot X^i \right)^2, \tag{7}$$

$$\sum_{j=-2\eta_2}^{2\eta_2} e_j X^j = \left(\sum_{i=-\eta_2}^{\eta_2} \frac{(2\eta_2)!}{(\eta_2 + i)!(\eta_2 - i)!} \cdot 2^{-2\eta_2} \cdot X^i \right)^2. \tag{8}$$

Proof. The probability is an intersection of two independent events, which can be written as:

$$P_{CBD} = \Pr(\Delta\mathbf{r} = 0) \cdot \Pr(\Delta\mathbf{e}_1 \in [\gamma_u]).$$

Let $\mathrm{P}_R(X), \mathrm{P}_{R'}(X), \mathrm{P}_E(X), \mathrm{P}_{E'}(X)$ be probability polynomials which represent the distributions of single coefficients in $\mathbf{r}$, $\mathbf{r}'$, $\mathbf{e}_1$, and $\mathbf{e}'_1$, respectively. By applying (2) in Sec. 2.3.2, the probability polynomials of single coefficient Δr and Δe_1 in $\Delta\mathbf{r}$ and $\Delta\mathbf{e}_1$ are given by:

$$\begin{aligned}
\mathrm{P}_{\Delta r}(X) &= \mathrm{P}_R(X) \cdot \mathrm{P}_{R'}(X) = \mathrm{P}_R(X)^2 \\
&= \left(\sum_{i=-\eta_1}^{\eta_1} \frac{(2\eta_1)!}{(\eta_1+i)!(\eta_1-i)!} \cdot 2^{-2\eta_1} \cdot X^i \right)^2 = \sum_{j=-2\eta_1}^{2\eta_1} r_j X^j \\
\mathrm{P}_{\Delta e_1}(X) &= \mathrm{P}_E(X) \cdot \mathrm{P}_{E'}(X) = \mathrm{P}_E(X)^2 \\
&= \left(\sum_{i=-\eta_2}^{\eta_2} \frac{(2\eta_2)!}{(\eta_2+i)!(\eta_2-i)!} \cdot 2^{-2\eta_2} \cdot X^i \right)^2 = \sum_{j=-2\eta_2}^{2\eta_2} e_j X^j.
\end{aligned}$$

Finally, the probability for a single coefficient should be raised to the power of nk to calculate the probability for all the coefficients in $\Delta\mathbf{r}$ and $\Delta\mathbf{e}_1$. Thus, we obtain the total probability of CBD injectivity as in (6), where r_0 is from (7) by setting $j = 0$ and e_j is from (8). □

Lemma 3 (MSIS Injectivity). *Let positive integers $q, k, n, d, \eta_1, \eta_2, d_u$ represent the parameters of Kyber (see Sec. 2.2), where for each distinct message pair $m \neq m'$ in the message space, we have $(\mathbf{r}, \mathbf{e}_1, e_2) := \mathsf{Sam}(G(m))$ and $(\mathbf{r}', \mathbf{e}'_1, e'_2) := \mathsf{Sam}(G(m'))$, where $\mathbf{r}, \mathbf{r}' \in R_q^k$ are sampled from $\beta_{\eta_1}^k$, $\mathbf{e}_1, \mathbf{e}'_1 \in R_q^k$ are sampled from $\beta_{\eta_2}^k$, and we define $\Delta\mathbf{r} := \mathbf{r} - \mathbf{r}'$ and $\Delta\mathbf{e}_1 := \mathbf{e}_1 - \mathbf{e}'_1$. The probability $P_{MSIS} := \Pr\left(\cup_{m\neq m'} \mathbf{A}^T\Delta\mathbf{r} + \Delta\mathbf{e}_1 \in [\gamma_u] \cap \Delta\mathbf{r} \neq \mathbf{0}\right)$ is upper bounded by:*

$$\begin{aligned}
P_{MSIS} \leq & \left(\frac{(4\eta_1+1)(4\eta_2+2\gamma_u+1)}{q} \right)^{nk} + \sum_{i=1}^{e} \binom{d}{i} \cdot \frac{1}{|W_i|^{2k}} \\
& \cdot \left\{ \frac{(4\eta_1 + 2\|W_i\|_\infty + 1) \cdot (4\eta_2 + 2\gamma_u + 2\|W_i\|_\infty + 1)}{q^{(1-i/d)}} \right\}^{nk}, \qquad (9)
\end{aligned}$$

where $\gamma_u := \left\lfloor \frac{q}{2^{d_u}} \right\rfloor$, e is the largest number such that $2\eta_1\sqrt{n} \geq q^{e/d}$, and W_i is a finite set defined in (10).

Proof. Let E denote the event $\cup_{m\neq m'} \mathbf{A}^T\Delta\mathbf{r} + \Delta\mathbf{e}_1 \in [\gamma_u] \cap \Delta\mathbf{r} \neq \mathbf{0}$. If E occurs, then there exist a pair of messages m, m' and corresponding vectors $\mathbf{r} \neq \mathbf{r}'$ (resp. $\mathbf{e}_1, \mathbf{e}'_1$ in the support of the distribution $\beta_{\eta_1}^k$ (resp. $\beta_{\eta_2}^k$), and a vector $\mathbf{e_u} \in [\gamma_u]$ such that

$$\mathbf{A}^T\Delta\mathbf{r} + (\Delta\mathbf{e}_1 - \mathbf{e_u}) = \mathbf{0}.$$

It follows that $P_{MSIS} := \Pr[E]$ is upper bounded by:

$$\begin{aligned}
&\Pr_{\mathbf{A}\leftarrow R_q^{k\times k}}\left[\exists(\Delta\mathbf{r}, \Delta\mathbf{e}_1 - \mathbf{e_u}) \in S_{\theta_1}^k \setminus \{\mathbf{0}\} \times S_{\theta_2}^k : \mathbf{A}^T\Delta\mathbf{r} + (\Delta\mathbf{e}_1 - \mathbf{e_u}) = \mathbf{0}\right] \\
&\quad\le \frac{|S_{\theta_1}|^k \cdot |S_{\theta_2}|^k}{q^{nk}} + \sum_{i=1}^{e} \frac{\binom{d}{i} \cdot \left|S_{\theta_1+\|W_i\|_\infty}\right|^k \cdot \left|S_{\theta_2+\|W_i\|_\infty}\right|^k}{|W_i|^{k+k} \cdot q^{nk(1-i/d)}} \\
&\quad= \left(\frac{|S_{\theta_1}| \cdot |S_{\theta_2}|}{q^n}\right)^k + \sum_{i=1}^{e} \binom{d}{i} \cdot \left(\frac{\left|S_{\theta_1+\|W_i\|_\infty}\right| \cdot \left|S_{\theta_2+\|W_i\|_\infty}\right|}{|W_i|^2 \cdot q^{n(1-i/d)}}\right)^k,
\end{aligned}$$

where in the first line, $\theta_1 \in \mathbb{N}$ (resp. $\theta_2 \in \mathbb{N}$) represent the maximum absolute value of coefficients of polynomials in the support set $S_{\theta_1}^k$ (resp. $S_{\theta_2}^k$) of $\Delta\mathbf{r}$ (resp. $\Delta\mathbf{e}_1 - \mathbf{e_u}$), and in the second line we apply Theorem 1 with $\boldsymbol{z_1}, \boldsymbol{z_2} = (\Delta\mathbf{r}, (\Delta\mathbf{e}_1 - \mathbf{e_u}))$. We first construct the sets S_{θ_1} and S_{θ_2}, then calculate the cardinality of the finite sets. For $\Delta\mathbf{r}$, since the support of the distribution β_{η_1} of $\mathbf{r}, \mathbf{r}'$ has maximal coefficient absolute value η_1, the maximal coefficient absolute value of $\Delta\mathbf{r}$ is $\theta_1 = 2\eta_1$. Similarly, for $\Delta\mathbf{e}_1 - \mathbf{e_u}$, since $\mathbf{e}_1, \mathbf{e}_1'$ are sampled from β_{η_2} and $\mathbf{e_u} \in [\gamma_u]$, the maximal coefficient absolute value $\theta_2 = 2\eta_2 + \gamma_u$.

Thus, we have:

$$\begin{aligned}
S_{\theta_1}^k = S_{2\eta_1}^k &:= \left\{\Delta\mathbf{r} \in R_q^k : \|\Delta\mathbf{r}\|_\infty \le 2\eta_1\right\}, \\
S_{\theta_2}^k = S_{2\eta_2+\gamma_u}^k &:= \left\{\Delta\mathbf{e}_1 - \mathbf{e_u} \in R_q^k : \|\Delta\mathbf{e}_1 - \mathbf{e_u}\|_\infty \le 2\eta_2 + \gamma_u\right\}.
\end{aligned}$$

Here, for $\alpha \in \mathbb{N}$, as defined in Theorem 1, $S_\alpha := \{y \in R_q : \|y\|_\infty \le \alpha\}$, and therefore the cardinality of S_α is:

$$|S_\alpha| = (2\alpha + 1)^n,$$

where n is the number of coefficients in a ring element. Thus, the first component in the probability bound above becomes:

$$\left(\frac{|S_{\theta_1}| \cdot |S_{\theta_2}|}{q^n}\right)^k = \left(\frac{|S_{2\eta_1}| \cdot |S_{2\eta_2+\gamma_u}|}{q^n}\right)^k = \left(\frac{(4\eta_1+1)(4\eta_2+2\gamma_u+1)}{q}\right)^{nk}.$$

The second component is dependent on the size of the set W_i. In the following, for each i, we define $t := \left\lfloor \frac{1}{2} q^{i/d} \right\rfloor$. We use the construction from Section 3.3 of [17]:

$$\begin{cases}
i = 1, & \begin{cases} |W_1| = 2n, \\ \|W_1\|_\infty = 1 \end{cases} \\
i \ge 2, & \begin{cases}
t < \sqrt{n}, & \begin{cases} |W_i| = \sum_{j=0}^{t^2} \binom{n}{j} \cdot 2^j, \\ \|W_i\|_\infty = 1 \end{cases} \\
t \ge \sqrt{n} & \begin{cases}
\textit{Set 1} & \begin{cases} |W_i| \ge V_n(\frac{1}{2} q^{i/d} - \sqrt{n}), \\ \|W_i\|_\infty = \left\lfloor \frac{1}{2} q^{i/d} \right\rfloor \end{cases} \\
\textit{Set 2} & \begin{cases} |W_i| = (2\left\lfloor \frac{t}{\sqrt{n}} \right\rfloor + 1)^n, \\ \|W_i\|_\infty = \left\lfloor \frac{t}{\sqrt{n}} \right\rfloor \end{cases}
\end{cases}
\end{cases}
\end{cases} \tag{10}$$

where $V_n(R)$ is the volume of n-dim ball with radius r, which can be calculated by $V_n(R) = \frac{(\pi/2)^{\lfloor \frac{n}{2} \rfloor}}{n!!}(2R)^n$. For $t \geq \sqrt{n}$, we choose the one which can produce a smaller bound between Set 1 and Set 2. Therefore, the second component is simplified as:

$$\begin{aligned}
&\sum_{i=1}^{e} \binom{d}{i} \cdot \left(\frac{\left|S_{\theta_1 + \|W_i\|_\infty}\right| \cdot \left|S_{\theta_2 + \|W_i\|_\infty}\right|}{|W_i|^2 \cdot q^{n(1-i/d)}} \right)^k \\
&= \sum_{i=1}^{e} \binom{d}{i} \cdot \left[\frac{\left|S_{2\eta_1 + \|W_i\|_\infty}\right| \cdot \left|S_{2\eta_2 + \gamma_u + \|W_i\|_\infty}\right|}{|W_i|^2 \cdot q^{n(1-i/d)}} \right]^k \\
&= \sum_{i=1}^{e} \binom{d}{i} \cdot \frac{1}{|W_i|^{2k}} \cdot \left\{ \frac{(4\eta_1 + 2\|W_i\|_\infty + 1) \cdot (4\eta_2 + 2\gamma_u + 2\|W_i\|_\infty + 1)}{q^{(1-i/d)}} \right\}^{nk}.
\end{aligned}$$

Now, we summarize the two components to generate our result for MSIS injectivity:

$$\begin{aligned}
&\Pr_{\mathbf{A} \leftarrow R_q^{k \times k}} \left[\exists (\Delta\mathbf{r}, \Delta\mathbf{e}_1 - \mathbf{e_u}) \in S_{\theta_1}^k \setminus \{\mathbf{0}\} \times S_{\theta_2}^k : \mathbf{A}^T \Delta\mathbf{r} + (\Delta\mathbf{e}_1 - \mathbf{e_u}) = \mathbf{0} \right] \\
&\quad \leq \left(\frac{(4\eta_1 + 1)(4\eta_2 + 2\gamma_u + 1)}{q} \right)^{nk} + \sum_{i=1}^{e} \binom{d}{i} \cdot \frac{1}{|W_i|^{2k}} \\
&\quad \cdot \left\{ \frac{(4\eta_1 + 2\|W_i\|_\infty + 1) \cdot (4\eta_2 + 2\gamma_u + 2\|W_i\|_\infty + 1)}{q^{(1-i/d)}} \right\}^{nk}.
\end{aligned}$$

This completes the proof. □

4 Numerical Result and Analysis

We now calculate the numerical bounds of injectivity of Kyber[1]. As mentioned

Table 1: Third Round Kyber Parameters from specification [3].

	n	d	q	k	η_1	η_2	(d_u, d_v)	δ	bit security (λ)
Kyber512	256	128	3329	2	3	2	(10, 4)	2^{-139}	128
Kyber768	256	128	3329	3	2	2	(10, 4)	2^{-164}	192
Kyber1024	256	128	3329	4	2	2	(11, 5)	2^{-174}	256

in Section 2, $n = 256$ is the fixed degree of the quotient polynomial, q is the module number of the module ring R_q, and d is the degree of splitting. The scheme

[1] The code can be accessed at https://github.com/RdWeirdo981/Injectivity-paper-codes.

Table 2: Kyber η-Injectivity

	M	P_{CBD}	P_{MSIS}	η	$\sqrt{\eta}$
Kyber512	2^{511}	2^{-1106}	2^{-354}	2^{-354}	2^{-177}
Kyber768	2^{511}	2^{-1445}	2^{-937}	2^{-934}	2^{-467}
Kyber1024	2^{511}	2^{-2420}	2^{-1687}	2^{-1687}	2^{-844}

uses k to represent dimension of secrete key vector, η_1 and η_2 to represent the sampling parameters for $\mathbf{s}, \mathbf{e}, \mathbf{r}$ and $\mathbf{e}_1, e_2$ respectively. From a security estimation perspective, δ is for δ-correctness, which is the probability of a decryption failure attack successfully happening. The security levels λ defined by Call for Proposals [22] of the three parameter sets are consistent with AES128 against 2^{170} MAXDEPTH quantum gates or 2^{143} classical gates as level 1, AES192 against 2^{233} MAXDEPTH quantum gates or 2^{207} classical gates as level 3, and AES256 against 2^{298} MAXDEPTH quantum gates or 2^{272} classical gates as level 5.

4.1 Numerical Analysis of η-injectivity

We first analyze the injectivity itself for concrete parameter sets. Overall, our result in Table 2 indicates the injectivity assumption in [13] holds for all parameter sets of CRYSTALS-Kyber by showing that $\sqrt{\eta} \leq 2^{-\lambda}$. We look further at Table 2 to analyze the separated components. Again, M, P_{CBD}, P_{SIS} are consistent with previous definitions, and η is calculated from $\eta \leq M \cdot P_{CBD} + P_{MSIS}$. The $\sqrt{\eta}$ is also calculated because we want to investigate the injectivity assumption in [13]. Some interesting observations are also summarized as below.

- We remark that the injectivity bounds obtained in this version of the paper improve upon the bounds in an earlier version [8], removing an extra unnecessary union bound for the computation P_{MSIS}, which is already accounted for in the bound of Theorem 1.
- Theorem 1 of [17] constructs 4 sets of W_i as summarized in (10). In the numerical calculation, for $t \geq \sqrt{n}$, Set 2 of W_i will always minimize the probability result rather than Set 1. This is because Set 1 uses the volume of n-dimensional ball with radius $\frac{1}{2}q^{i/d} - \sqrt{n}$. If we take a further look at the equation of the volume, it can be calculated that the constant part $\frac{(\pi/2)^{\lfloor n/2 \rfloor}}{n!!}$ is really big, which makes Set 1 always to generate the larger result than Set 2. When the iteration of the sum goes up to around 65, we will obtain the largest single round result as seen in Fig. 1a, Fig. 1c, and Fig. 1e. And that is why the final cumulative result goes large in Fig. 1b, Fig. 1d, and Fig. 1f.
- We also observed that Kyber512 has an apparent discontinuous point after round 65. This is due to $e = \left\lfloor d \cdot \log_q(2\eta_1\sqrt{n}) \right\rfloor$. Kyber512 has $\eta_1 = 2$ rather than 3 in Kyber768 and Kyber1024.

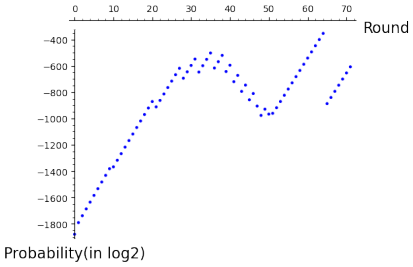

(a) Single Result for Kyber512

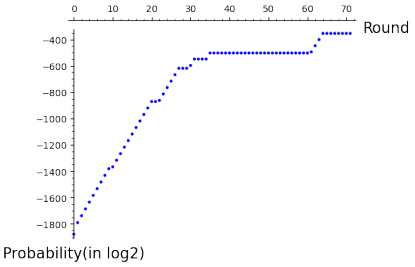

(b) Cumulative Result for Kyber512

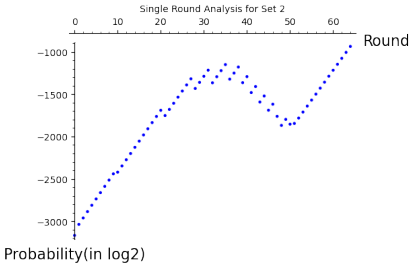

(c) Single Result for Kyber768

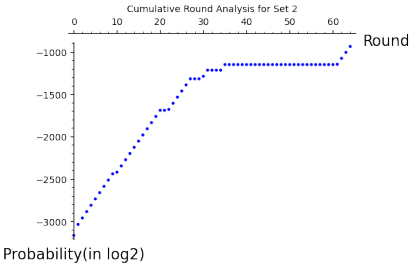

(d) Cumulative Result for Kyber768

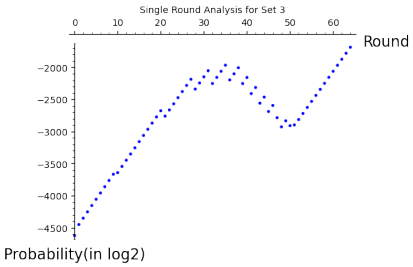

(e) Single Result for Kyber1024

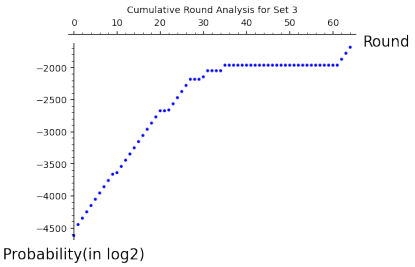

(f) Cumulative Result for Kyber1024

Fig. 1: Output of τ Function in MSIS Part

4.2 Implications of the relationship between η-injectivity and QROM advantage

We now analyze the implications of η-injectivity in Corollary 4.7 in [13]. We first recall the latter Corollary below. Here, given a randomized public key encryption scheme P and a random oracle G, we denote by $\mathsf{P}' = T(\mathsf{P}, G)$ the deterministic public key encryption scheme obtained from P by deriving the encryption scheme randomness by applying G to the input message. We denote by $\mathsf{FO}^{\not\perp}(\mathsf{P}, \mathsf{F}, G, H) = U^{\not\perp}(\mathsf{P}', \mathsf{F}, H)$ the KEM obtained by applying the $\mathsf{U}^{\not\perp}$ Fujisaki-Okamoto (FO) transform, an implicit rejection variant of FO transform defined in [12].

Corollary 1 ([13], Cor. 4.7). *Let P be a δ-correct randomized public key encryption scheme with message space $\mathcal{M}$ and randomness space $\mathcal{R}$. Let G and H be quantum accessible random oracles, and F be a PRF. Suppose that $\mathsf{P}' = T(\mathsf{P}, G)$ is η-injective and let $\mathsf{FO}^{\not\perp}(\mathsf{P}, \mathsf{F}, G, H) = U^{\not\perp}(\mathsf{P}', \mathsf{F}, H)$. Let $\mathcal{A}$ be an adversary with run-time $T_{\mathcal{A}}$ against the $\mathsf{IND\text{-}CCA}$ security of $\mathsf{FO}^{\not\perp}(\mathsf{P}, \mathsf{F}, G, H)$ issuing at most q_G (resp. q_H) quantum queries to G (resp. H) with query depth at most d_G (resp. d_H) and at most q_{dec} classical queries to the decapsulation oracle of $\mathsf{FO}^{\not\perp}(\mathsf{P}, \mathsf{F}, G, H)$. Then, we can construct an $\mathsf{IND\text{-}CPA}$ adversary $\mathcal{B}_1$ against P and a PRF adversary $\mathcal{B}_2$ against F issuing at most q_{dec} queries, satisfying:*

$$\begin{aligned}\mathsf{Adv}^{\mathsf{IND-CCA}}_{\mathsf{FO}^{\not\perp}(\mathsf{P},\mathsf{F},G,H)}(\mathcal{A}) \leq\ & 8d_H \cdot (d_G+1) \cdot \left(\mathsf{Adv}^{\mathsf{IND-CPA}}_{\mathsf{P}}(\mathcal{B}_1) + \frac{8\cdot(3q_G+1)}{|\mathcal{M}|}\right) \\ & + 6\cdot(3q_G+q_{dec}) \cdot \left((8d_G+1)\cdot\delta + \sqrt{3\eta}\right) \\ & + (4d_H+12)\cdot\eta + 2\mathsf{Adv}^{\mathsf{PRF}}_{\mathsf{F}}(\mathcal{B}_2)\,. \end{aligned} \tag{11}$$

We note that above, for the QROM queries to G and H, d_G (resp. d_H) is the algorithm's query depth to G (resp. H), consisting of d_G (resp. d_H) sequential bunches of n_G (resp. n_H) parallelized queries to G (resp. H), where n_G (resp n_H) is the parallelization factor, and $q_G = n_G \cdot d_G$ (resp. $q_H = n_H \cdot d_H$) is the total number of oracle queries to G (resp. H).

In the following analysis of Kyber, we omit the parameter $h = \mathsf{H}(pk)$ that is not needed in the Kyber IND-CCA security reduction (as also noted in [3]), and is only included for robustness against multi-target attacks. Indeed, note that pk is fixed anyway throughout the attack for the standard (single target key) IND-CCA security model, so it does not affect the single target key IND-CCA security model under the scope in this paper.

To apply Corollary 1 (Corollary 4.7 of [13]) for $\mathsf{FO}^{\not\perp}(\mathsf{P}, \mathsf{F}', G', H') = \mathsf{U}^{\not\perp}(T(\mathsf{P}, G'), \mathsf{F}', H')$ to KYBER.$\mathsf{CCAKEM}(\mathsf{KDF}, G, H)$ specified in [3], we observe that KYBER.$\mathsf{CCAKEM}(\mathsf{KDF}, G, H)$ can be viewed as $\mathsf{FO}^{\not\perp}(\mathsf{P}, \mathsf{F}', G', H')$ with:

$$\begin{aligned} G'(m) &:= G_2(m) \\ H'(m,c) &:= \mathsf{KDF}(G_1(m), H(c)) \\ \mathsf{F}'(z,c) &:= \mathsf{KDF}(z, H(c)) \end{aligned} \tag{12}$$

Here, $G(m) := (G_1(m), G_2(m))$ is split into two sub-functions G_1, G_2 that compute the first and second parts of the output of G in KYBER.CCAKEM. In the following, we apply Corollary 1 combined with our injectivity results to analyze the concrete IND-CCA security of the simplified $\mathsf{FO}^{\not\perp}$ model $\widehat{\text{KYBER}}.\mathsf{CCAKEM}(\mathsf{P}, \mathsf{F}', G', H') := \mathsf{FO}^{\not\perp}(\mathsf{P}, \mathsf{F}', G', H')$ (where P denotes the Kyber.CPA scheme) of KYBER.$\mathsf{CCAKEM}(\mathsf{KDF}, G, H)$. In our simplified model $\widehat{\text{KYBER}}.\mathsf{CCAKEM}(\mathsf{P}, \mathsf{F}', G', H')$ of KYBER.$\mathsf{CCAKEM}(\mathsf{KDF}, G, H)$, H', F' are modelled as a fresh quantum random oracle (resp. PRF), rather than being implemented from underlying random oracles KDF, G_1, H using 'double hashing' as in the right-hand side of (12). We remark that the security of the actual 'double-hashing' variant of KYBER.CCAKEM can be related to the security of the 'single hashing' model $\widehat{\text{KYBER}}.\mathsf{CCAKEM}$ we discuss in this Section, either via indifferentiability arguments in the QROM [26] or via tighter direct reductions, as analyzed in the very recent work of Maram and Xagawa [16]; however we do not analyze this relation concretely in this Section.

To summarize, for our simplified $\mathsf{FO}^{\not\perp}$ model $\widehat{\text{KYBER}}.\mathsf{CCAKEM}(\mathsf{P}, \mathsf{F}', G', H')$ of KYBER.CCAKEM, the IND-CCA attacker advantage bound has the following form:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{IND-CCA}}_{\mathsf{FO}^{\not\perp}(\mathsf{P},\mathsf{F}',G',H')}(\mathcal{A}) \leq\ & 8d_{H'} \cdot (d_{G'}+1) \cdot \left(\mathsf{Adv}^{\mathsf{IND-CPA}}_{\mathsf{P}}(\mathcal{B}_1) + \frac{8 \cdot (3q_{G'}+1)}{2^{256}} \right) \\ & + 6 \cdot (3q_{G'} + q_{dec}) \cdot \left((8d_{G'}+1) \cdot \delta + \sqrt{3\eta} \right) \\ & + (4d_{H'}+12) \cdot \eta + 2\mathsf{Adv}^{\mathsf{PRF}}_{\mathsf{F}'}(\mathcal{B}_2)\,, \end{aligned} \tag{13}$$

where concrete bounds for δ and η for Kyber512, Kyber768 and Kyber1024 are given in Table 1 and Table 2, respectively.

Note that under the pseudorandomness of G and PRF in the standard model and the QROM for XOF (that generates the public $\mathbf{A}$ matrix from a seed), there is a straightforward tight reduction from MLWE to IND-CPA security of KYBER.CCAKEM based on the pseudorandomness of the public key and ciphertext as in Theorem 1 in [3], with a multiplicative loss factor of 2. Therefore one can replace $\mathsf{Adv}^{\mathsf{IND-CPA}}_{\text{KYBER.CPAPKE}}(\mathcal{B}_1)$ by $\approx 2\mathsf{Adv}^{\text{MLWE}}(\mathcal{B'}_1)$ for an attacker $\mathcal{B'}_1$ related to $\mathcal{B}_1$. We also observe that there is no square-root loss in the advantage, and for highly parallelized quantum attacks, with $d_{G'}, d_{H'} = O(1)$, the result of (13) is nearly tight. Therefore, against such attacks, there is a nearly tight QROM security proof for Kyber's CCA security from MLWE.

Finally, as an illustration, we give a concrete interpretation and example rough estimates of the reduction loss parameters in (13) for highly parallelised attacks on Kyber with its realization of random oracles G', H' with concrete hash functions. We first note that the total quantum attack circuit depth, called MAXDEPTH in the NIST PQC process [22], is related to parameters in (13) as:

$$\text{MAXDEPTH} = d_{\mathcal{A}} + d_{G'} \cdot \text{DEPTH}_{G'} + d_{H'} \cdot \text{DEPTH}_{H'},$$

where $\text{DEPTH}_{G'}$ (resp. $\text{DEPTH}_{H'}$) is the depth of the quantum circuit implementing G' (resp. H'), and $d_{\mathcal{A}}$ is the depth of $\mathcal{A}$ excluding the cir-

cuits implementing the random oracle realization. For Kyber, $\mathrm{DEPTH}_{G'}$ and $\mathrm{DEPTH}_{H'}$ are the depth of a quantum circuit for SHA3-512 and max(DEPTH(SHA3-512, SHA3-256)) + DEPTH(SHAKE-256)), respectively. Given concrete values for the quantum circuit depth of the hash functions and total attack MAXDEPTH, one can get concrete estimates for $d_{G'}$ and $d_{H'}$. For example, MAXDEPTH $= 2^{20}$ and $\mathrm{DEPTH}_{G'}, \mathrm{DEPTH}_{H'} \approx 2^{13}$ (the latter approximate depth estimates are based on the quantum circuit depth estimate for SHA3-256 in Table 2 of [2]) implies $d_{G'}, d_{H'} \leq 2^7$. In this example, we have the following concrete bound for IND-CCA security of Kyber512:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{IND-CCA}}_{\mathsf{FO}^{\not\perp}(\mathsf{P},\mathsf{F}',G',H')}(\mathcal{A}) \leq\ & 2^{17} \cdot \left(\mathsf{Adv}^{\mathsf{IND-CPA}}_{\mathrm{KYBER.CPAPKE}}(\mathcal{B}_1) + \frac{q_{G'}}{2^{251}}\right) \\ & + (9q_{G'} + 3q_{dec})\left(2^{-128} + 2^{-176}\right) \\ & + 2^{-345} + 2\mathsf{Adv}^{\mathsf{PRF}}_{\mathsf{F}'}(\mathcal{B}_2)\,.\end{aligned}$$

Similarly, for Kyber768 and Kyber1024, respectively:

$$\begin{aligned}\mathsf{Adv}^{\mathsf{IND-CCA}}_{\mathsf{FO}^{\not\perp}(\mathsf{P},\mathsf{F}',G',H')}(\mathcal{A}) \leq\ & 2^{17} \cdot \left(\mathsf{Adv}^{\mathsf{IND-CPA}}_{\mathrm{KYBER.CPAPKE}}(\mathcal{B}_1) + \frac{q_{G'}}{2^{251}}\right) \\ & + (9q_{G'} + 3q_{dec})\left(2^{-153} + 2^{-466}\right) \\ & + 2^{-925} + 2\mathsf{Adv}^{\mathsf{PRF}}_{\mathsf{F}'}(\mathcal{B}_2)\,,\end{aligned}$$

$$\begin{aligned}\mathsf{Adv}^{\mathsf{IND-CCA}}_{\mathsf{FO}^{\not\perp}(\mathsf{P},\mathsf{F}',G',H')}(\mathcal{A}) \leq\ & 2^{17} \cdot \left(\mathsf{Adv}^{\mathsf{IND-CPA}}_{\mathrm{KYBER.CPAPKE}}(\mathcal{B}_1) + \frac{q_{G'}}{2^{251}}\right) \\ & + (9q_{G'} + 3q_{dec})\left(2^{-163} + 2^{-843}\right) \\ & + 2^{-1678} + 2\mathsf{Adv}^{\mathsf{PRF}}_{\mathsf{F}'}(\mathcal{B}_2)\,.\end{aligned}$$

Note that the above rough estimates may not be very accurate but are included to give some idea of how our result could be interpreted given concrete assumptions on attack MAXDEPTH and hash function realization circuit depth. The probability of decryption failure error δ also plays a role on the RHS, and in fact, as can be seen above for the Kyber parameter choices, it becomes the dominant term compared to the negligible injectivity terms (to clarify the contributions of decryption error vs. injectivity, we have shown both terms in the above bounds; the first term in the factor multiplying $(9q_{G'} + 3q_{dec})$ is the contribution of the decryption error δ and the second term is due to the injectivity). We refer the readers to Sec. 4 and Sec. 5 in Kyber's Specification document [3] for more details on the choice of δ in Kyber.

5 Conclusion

We have followed the work of [13] and taken a step further to investigate and analyze the injectivity of CRYSTALS-Kyber, which uses FO transform to convert an IND-CPA public key encryption into an IND-CCA key encapsulation mechanism. The theoretical bound and corresponding numerical results are provided and have shown that all the parameter sets are reasonably safe with respect to

the injectivity assumption in [5, 13], which means that the collision probability of two different messages having the same output ciphertext is negligible in these parameter sets. This allows us to obtain the tightest QROM CCA security bounds to date for the 'single hashing' model of Kyber CCAKEM against attacks with low quantum circuit depth, for the Kyber parameter sets.

Other schemes using the FO transform like Saber, can be investigated by a similar method. One main reason why we cannot immediately apply the steps for calculating Kyber injectivity for Saber is that it uses a power-of-2 rather than a prime modulus. It makes the approach in [17] not directly applicable for Saber.

Acknowledgement. This work was supported in part by Australian Research Council Discovery Grant DP180102199.

References

1. Ambainis, A., Hamburg, M., Unruh, D.: Quantum security proofs using semi-classical oracles. In: Advances in Cryptology - CRYPTO 2019. pp. 269–295. Springer (2019). https://doi.org/10.1007/978-3-030-26951-7_10
2. Amy, M., Matteo, O.D., Gheorghiu, V., Mosca, M., Parent, A., Schanck, J.M.: Estimating the cost of generic quantum pre-image attacks on SHA-2 and SHA-3. In: Selected Areas in Cryptography - SAC 2016. LNCS, vol. 10532, pp. 317–337. Springer (2016). https://doi.org/10.1007/978-3-319-69453-5_18
3. Avanzi, R., Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J., Schwabe, P., Seiler, G., Stehlé, D.: Algorithm specifications and supporting documentation (version 3.02). Tech. rep., Submission to the NIST post-quantum project (2021), available at https://pq-crystals.org/kyber/resources.shtml
4. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Proceedings of the 1st ACM Conference on Computer and Communications Security. pp. 62–73. Association for Computing Machinery (1993). https://doi.org/10.1145/168588.168596
5. Bindel, N., Hamburg, M., Hövelmanns, K., Hülsing, A., Persichetti, E.: Tighter proofs of cca security in the quantum random oracle model. In: Theory of Cryptography. pp. 61–90. Springer (2019). https://doi.org/10.1007/978-3-030-36033-7_3
6. Boneh, D., Dagdelen, Ö., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random oracles in a quantum world. In: Advances in Cryptology - ASIACRYPT 2011. pp. 41–69. Springer (2011). https://doi.org/10.1007/978-3-642-25385-0_3
7. Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehlé, D.: Crystals-kyber: a cca-secure module-lattice-based kem. In: 2018 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 353–367. IEEE (2018). https://doi.org/10.1109/EuroSP.2018.00032
8. Ding, X., Esgin, M.F., Sakzad, A., Steinfeld, R.: An injectivity analysis of crystals-kyber and implications on quantum security. In: ACISP. Lecture Notes in Computer Science, vol. 13494, pp. 332–351. Springer (2022). https://doi.org/10.1007/978-3-031-22301-3_17

9. Fujisaki, E., Okamoto, T.: How to enhance the security of public-key encryption at minimum cost. In: International Workshop on Public Key Cryptography. pp. 53–68. Springer (1999). https://doi.org/10.1007/3-540-49162-7_5
10. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Advances in Cryptology - CRYPTO 1999. pp. 537–554. Springer (1999). https://doi.org/10.1007/3-540-48405-1_34
11. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the twenty-eighth annual ACM symposium on Theory of Computing. pp. 212–219. STOC '96, Association for Computing Machinery (1996). https://doi.org/10.1145/237814.237866
12. Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the fujisaki-okamoto transformation. In: Theory of Cryptography. Lecture Notes in Computer Science, vol. 10677, pp. 341–371. Springer (2017). https://doi.org/10.1007/978-3-319-70500-2_12
13. Kuchta, V., Sakzad, A., Stehlé, D., Steinfeld, R., Sun, S.F.: Measure-rewind-measure: Tighter quantum random oracle model proofs for one-way to hiding and cca security. In: Advances in Cryptology - EUROCRYPT 2020. pp. 703–728. Springer (2020). https://doi.org/10.1007/978-3-030-45727-3_24
14. Lily, C.N., Stephen, J.N., Yi-Kai, L.N., Dustin, M.N., Rene, P.N., Ray, P.N., Daniel, S.T.N.: Report on post-quantum cryptography. Tech. rep., National Institute of Standards and Technology, Gaithersburg, MD (2016). https://doi.org/10.6028/NIST.IR.8105
15. Lyubashevsky, V.: Basic lattice cryptography: Encryption and fiat-shamir signatures (2019)
16. Maram, V., Xagawa, K.: Post-quantum anonymity of kyber. Cryptology ePrint Archive, Paper 2022/1696 (2022), available at https://eprint.iacr.org/2022/1696
17. Nguyen, N.K.: On the non-existence of short vectors in random module lattices. In: Advances in Cryptology - ASIACRYPT 2019. pp. 121–150. Springer (2019). https://doi.org/10.1007/978-3-030-34621-8_5
18. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM review **41**(2), 303–332 (1999). https://doi.org/10.1137/S0097539795293172
19. Stehlé, D., Steinfeld, R.: Making NTRU as secure as worst-case problems over ideal lattices. In: EUROCRYPT. Lecture Notes in Computer Science, vol. 6632, pp. 27–47. Springer (2011). https://doi.org/10.1007/978-3-642-20465-4_4
20. Stehlé, D., Steinfeld, R.: Making ntruencrypt and ntrusign as secure as standard worst-case problems over ideal lattices. Cryptology ePrint Archive, Paper 2013/004 (2013), available at https://eprint.iacr.org/2013/004
21. Sullivan, N.: Securing the post-quantum world (Sep 2021), available at https://blog.cloudflare.com/securing-the-post-quantum-world/
22. The NIST PQC Team: Submission requirements and evaluation criteria for the post-quantum cryptography standardization process (2017), available at https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf
23. The NIST PQC Team: PQC standardization process: Announcing four candidates to be standardized, plus fourth round candidates (2022), available at https://csrc.nist.gov/News/2022/pqc-candidates-to-be-standardized-and-round-4
24. Unruh, D.: Revocable quantum timed-release encryption. Journal of the ACM (JACM) **62**(6), 1–76 (2015). https://doi.org/10.1145/2817206

25. Weibel, A.: Round 2 post-quantum TLS is now supported in aws kms (Nov 2020), available at https://aws.amazon.com/blogs/security/round-2-post-quantum-tls-is-now-supported-in-aws-kms/
26. Zhandry, M.: How to record quantum queries, and applications to quantum indifferentiability. In: CRYPTO (2). Lecture Notes in Computer Science, vol. 11693, pp. 239–268. Springer (2019). https://doi.org/10.1007/978-3-030-26951-7_9