

Faster Lattice-Based KEMs via a Generic Fujisaki-Okamoto Transform Using Prefix Hashing

Julien Duman 
Ruhr Universität Bochum
Germany
julien.duman@rub.de

Kathrin Hövelmanns 
TU Eindhoven
The Netherlands
kathrin@hoevelmanns.de

Eike Kiltz 
Ruhr Universität Bochum
Germany
eike.kiltz@rub.de

Vadim Lyubashevsky
IBM Research Europe, Zurich, Switzerland
vad@zurich.ibm.com

Gregor Seiler
ETH Zürich and IBM Research Europe, Zurich
Switzerland
gseiler@inf.ethz.ch

ABSTRACT

Constructing an efficient CCA-secure KEM is generally done by first constructing a passively-secure PKE scheme, and then applying the Fujisaki-Okamoto (FO) transformation. The original FO transformation was designed to offer security in a single user setting. A stronger notion, known as multi-user security, considers the attacker’s advantage in breaking one of many user’s ciphertexts. Bellare et al. (EUROCRYPT 2020) showed that standard single user security implies multi-user security with a multiplicative tightness gap equivalent to the number of users.

To obtain even more confidence in the security of KEMs in the multi-user setting, it is a common design paradigm to also “domain separate” the random oracles of each user by including his public key as an input to the hash function. We are not aware of any formal analysis of this technique, but it was at least informally thought to be a computationally cheap way to add security. This design principle was carried over into the FO transformations used by several schemes in the NIST post-quantum standardization effort – notably the lattice-based schemes Kyber and Saber, which are two of the four KEM finalists.

In this work, we formally analyze domain separation in the context of the FO transformation in the multi-user setting. We first show that including the public key in the hash function is indeed important for the tightness of the security reductions in the ROM and the QROM. At the same time, we show that including the *entire* public key into the hash function is unnecessarily wasteful – it is enough to include just a small (e.g. 32 byte) unpredictable part of the key to achieve the same security. Reducing the input of the hash function results in a very noticeable improvement in the running time of the lattice-based KEMs. In particular, using this generic transform results in a 2X - 3X speed-up over the current (Round 3) key generation and encapsulation procedures in Kyber, and up to a 40% improvement in the same functions in Saber.

KEYWORDS

FO Transform; QROM; Key Exchange; Lattice Cryptography; Implementation

1 INTRODUCTION

Security definitions for public key encryption (PKE) schemes are generally stated in the *single-user* setting. In this setting, one party

publishes its public key, which allows other parties to send it encrypted messages of their choice. For practical applications, however, this was shown to not be enough. In particular, Håstad showed [20] that if there are multiple receivers, each with a different public key, and a sender encrypts the same message to all of them, then for certain RSA parameter settings everyone will be able to recover the message.

Håstad’s simple attack against the basic RSA cryptosystem demonstrated that schemes can be secure in the single user setting, but completely broken in the multi-user one. In the multi-user and multi-challenge (multi-user/challenge) setting, it is furthermore required that the encryptions of many plaintexts, possibly encrypted to different public keys, remain hidden secure. Furthermore, NIST lists resistance to multi-key attacks as a “desirable property” in their Call for Proposals for the post-quantum standardization process [31, Section 4.A.6]. Luckily, Bellare et al. [6] showed that for CPA and CCA-secure schemes, security in a single-user and single-challenge setting implies security (with a multiplicative loss in the number of parties times the number of challenges) in the multi-user/challenge setting as well.

Constructing an efficient Key Encapsulation Mechanism (KEM) with security against chosen-ciphertext attacks (CCA) is generally done by first constructing a passively-secure PKE scheme, and then applying the Fujisaki-Okamoto (FO) transformation [16, 17, 21]. To achieve multi-user security in the random oracle model, it is a common design paradigm to use “domain separation” so that in the random oracle model, the parties all appear to be using different random functions. The simple way of achieving this in the context of the FO transformation is to always include the public key as an additional input to the cryptographic hash function that is being modeled as a random oracle. The other stated reason for including the public key in the hash is an informally-defined notion of making the KEM “contributory” – that is, both parties affect the shared key.

For classical schemes based on the hardness of the discrete logarithm problem, including the public key as an extra input to the hash function has virtually no effect on the running time of the full protocol. The reason is that the key is short (e.g. 32 bytes), and the hash function contributes a negligible amount of computation compared to the much more expensive group operations such as exponentiation. So even if unnecessary for practical security, adding the public key into the hash function does not have any measurable negative effects on the scheme.

The above-mentioned domain separation technique was carried over to schemes in the ongoing NIST post-quantum standardization process. For example, Kyber [12] and Saber [14], which are two of the four KEM finalists in the NIST quantum-safe competition, use this method with the explicit purpose of protecting against multi-user attacks. Even though it was at least informally thought to be more secure, we are not aware of any formal security analysis of it.¹

1.1 Our Results

In this paper we formally analyze domain separation in the context of the FO transformation. Our results are twofold. First, we observe that there are good reasons for including the public key into the hash. We show that hashing the public key results in a tighter reduction in a multi-user/multi-challenge setting when converting a CPA-secure scheme into a CCA-secure one, than if one were to directly apply the hybrid argument from [6] to the FO transformation. And for schemes that additionally have a small correctness error, the tightness in the reduction is potentially even more improved. We additionally give a proof in the QROM which is significantly tighter than what one would trivially obtain from the hybrid approach of [6]. This proof also appears to use (but in a different way) the fact that each party uses a different random function.

Our second, and main, result is that even though there are reasons to hash the public key, hashing the *entire* public key (as is currently done) is unnecessarily wasteful. The sizes of public keys in lattice-based schemes (≈ 1 KB) are noticeably larger than the 32 byte keys used in the discrete log setting; so the hash function that takes the public key as input is now approximately an order of magnitude slower. At the same time, the underlying lattice-based CPA-secure scheme is significantly *more* computationally efficient than its discrete logarithm counterpart. When compounded, these two properties result in the hash function being a very significant contributor to the total running time of the resulting CCA-secure scheme.

FO Transformation with Prefix Hashing. Our new variant of the FO transformation, $\text{FO}_{\text{ID}(pk),m}^{\perp}$, transforms a passively secure PKE scheme into an actively secure KEM. It uses “implicit rejection” and “prefix hashing”: Implicit rejection (indicated by the \perp symbol in $\text{FO}_{\text{ID}(pk),m}^{\perp}$) means that decapsulation of an invalid ciphertext deterministically outputs a pseudorandom key; prefix hashing (indicated by the $\text{ID}(id),m$ in $\text{FO}_{\text{ID}(pk),m}^{\perp}$) means that we feed a small (e.g. 32 byte), unpredictable part of the public key into the hash function. FO_m^{\perp} also uses implicit rejection but does not include any part of pk in the hash function. Compared to feeding the whole public key into the hash function, our new transformation significantly reduces the running time of the scheme. In the ROM we prove that multi-user/challenge CPA-security of PKE tightly implies multi-user/challenge CCA-secure KEM. As an additional result, we give a reduction for multi-user/challenge security in the QROM which is tighter than the previously-known results which stem from the generic Bellare et al. result [6]. We remark that ThreeBears [18] (a Round 2 candidate of the NIST post-quantum competition) implicitly uses an FO transformation with prefix hashing and proves

¹This is in contrast to digital signatures where domain separation provably helps improving security in the multi-user setting. See, e.g., [7, 9].

multi-user/challenge CCA-security in the QROM [19]. It is unclear, however, if the specific proof carries over to the general FO transformation.

Impact on Kyber. Instantiating our transformation $\text{FO}_{\text{ID}(pk),m}^{\perp}$ with the CPA-secure Kyber PKE scheme, the key generation and encryption of Kyber is reduced by 30-56% and 47-66%, respectively; and the security of the CCA-scheme in the ROM and QROM is improved as in Fig. 1. So we are now in the same situation with lattice schemes as we are in the discrete logarithm setting. There is no longer a theory vs. practice trade-off required for achieving multi-user/challenge security – hashing the (partial) public key leads to tighter security reductions and is computationally very cheap to implement. Since Kyber is currently the most efficient – as measured by the running time of key generation, encapsulation, and decapsulation procedures [8] – of all KEMs (quantum-safe and non-quantum-safe), our new transformation also improves on how fast a secure KEM can be.

1.2 Impact on Concrete Security

In the Fujisaki-Okamoto transform, a hash function F (modeled as a random oracle) is mainly used to derive the PKE randomness r and the symmetric KEM key K from a random message m . In the following table, we define three variants of the “implicit rejection” Fujisaki-Okamoto transformation, depending on which parts of pk it includes in the hash function F .

Transformation	Use of hash F
FO_m^{\perp} [21]	$(K, r) = F(m)$
$\text{FO}_{\text{ID}(pk),m}^{\perp}$ (§3)	$(K, r) = F(\text{ID}(pk), m)$
$\text{FO}_{pk,m}^{\perp}$ (§3)	$(K, r) = F(pk, m)$

The formal definitions of $\text{FO}_{\text{ID}(pk),m}^{\perp}$ and $\text{FO}_{pk,m}^{\perp}$ are given in Fig. 4 of Section 3. FO_m^{\perp} was formally analyzed in [21]. We remark that [21] also considered a variant (called $\text{FO}_{m,c}^{\perp}$) where the ciphertext c is included in F . Follow-up work [10] showed that hashing c is not necessary so we do not further consider it here.

Fig. 1 compares multi-user/challenge security bounds in the ROM/QROM of FO_m^{\perp} , $\text{FO}_{\text{ID}(pk),m}^{\perp}$, and $\text{FO}_{pk,m}^{\perp}$. The ROM bounds for FO_m^{\perp} are new bounds obtained in this work. They improve on the naive “hybrid bound” $\text{Adv}_{\text{KEM}}^{(n,q_C)\text{-IND-CCA}} \leq nq_C \cdot \text{Adv}_{\text{KEM}}^{(n,q_C)\text{-IND-CCA}} \leq nq_C \cdot (\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} + q_F \delta(1))$ obtained by applying the Bellare et al. hybrid argument to the single user/challenge CCA-bounds of [21]. Unfortunately, we were not able to improve the hybrid bounds in the QROM setting for FO_m^{\perp} . The bounds for $\text{FO}_{\text{ID}(pk),m}^{\perp}$ and $\text{FO}_{pk,m}^{\perp}$ are new bounds obtained in this work, again improving on the hybrid bounds. Our proofs in the ROM rely on techniques of [21], while the ones in the QROM also rely on techniques of [10, 23, 26].

First we note that the bounds for $\text{FO}_{\text{ID}(pk),m}^{\perp}$ and $\text{FO}_{pk,m}^{\perp}$ are exactly the same (for prefixes with sufficient min-entropy l). Hence from a security perspective, it does not seem to make much sense to include the *entire* public key in the hash function. Hence in what follows we will only consider the prefix-hash variant $\text{FO}_{\text{ID}(pk),m}^{\perp}$. How the bounds for FO_m^{\perp} and $\text{FO}_{\text{ID}(pk),m}^{\perp}$ compare to each other, depends on the relation between $\delta(n)$ and δ and $\text{Adv}_{\text{PKE}}^{(n,q_C)\text{-IND-CPA}}$ and $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}}$, respectively.

Figure 1: Multi-user/challenge security advantages $\text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}$ in the ROM/QROM (simplified) for the Fujisaki-Okamoto variants FO_m^\perp , $\text{FO}_{\text{ID}(pk), m}^\perp$, and $\text{FO}_{pk, m}^\perp$, as described in the text. Here $\delta(n)$ is the n -user correctness error of PKE, q_C is the number of challenges, q_F is the number of (Q)ROM queries, and l is the min-entropy of $\text{ID}(pk)$.

FO variant	$\text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}$ (ROM)	$\text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}$ (QROM)
$\text{FO}_{\text{ID}(pk), m}^\perp$ (Th. 3.1+3.2)	$\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}} + q_F \delta(n) + \frac{n^2}{2^\ell}$	$\sqrt{q_F \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}} + q_F^2 \delta(n) + \frac{n^2}{2^\ell}}$
$\text{FO}_{pk, m}^\perp$ (Th. 3.1+3.2)	$\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}} + q_F \delta(n)$	$\sqrt{q_F \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}} + q_F^2 \delta(n)}$
FO_m^\perp (Th. 3.1+[6, 21])	$\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}} + n \cdot q_F \delta(1)$	$n q_C \cdot (\sqrt{q_F \text{Adv}_{\text{PKE}}^{\text{IND-CPA}}} + q_F \delta(1))$

The n -user correctness error $\delta(n)$. $\delta(n)$ is defined as the n -user correctness error, i.e., the probability that one message induces a decryption error for one of n independent public-keys. Informally (see Section 2.1 for a formal definition), the adversary, who is given n public and secret keys, chooses a message and one of the public keys under which the message is to be encrypted; he wins if the ciphertext then induces a decryption failure. The trivial bounds are $\delta(1) \leq \delta(n) \leq n\delta(1)$.

There is also very strong evidence that for most natural lattice-based schemes, $\delta(n) < n\delta(1)$. The decryption error of a basic LWE-based encryption scheme encrypting one bit with secret \vec{s} and encryption randomness \vec{r} is

$$\Pr_{\vec{s}, \vec{r} \leftarrow \psi^m} [\|\vec{s} \cdot \vec{r}\| > \lfloor q/4 \rfloor - 1].$$

If the coefficients of randomness \vec{r} were continuous Gaussians, rather than e.g. discrete binomials, then one could prove that the optimal strategy for the adversary is to “attack” the public key whose corresponding secret key has the largest norm. Then the probability that we are interested in is

$$\Pr_{\vec{s}_1, \dots, \vec{s}_n, \vec{r} \leftarrow \psi^m} [\|\vec{s}_i \cdot \vec{r}\| > \lfloor q/4 \rfloor - 1 : i := \arg \max_{1 \leq j \leq n} \|\vec{s}_j\|].$$

In schemes where the randomness is discrete, this is almost certainly still the best strategy. So, for example, in schemes where all the secret keys are prescribed to have the same norm, it makes no difference which public key to attack, and therefore we would have $\delta(n) \approx \delta(1)$. In Kyber and Saber, where each coefficient is chosen independently, there will be secret keys with larger norms, and so $\delta(n) > \delta(1)$. Nevertheless, because the norm of the secret keys has a tight concentration when the coefficients follow the binomial distribution, we will still have $\delta(n) < n \cdot \delta(1)$.

In Table 1, we give the exact values for $\delta(n)$ illustrating this phenomenon for an LWE scheme that has secret vectors consisting of $-1/0/1$ coefficients. For example, note that for $n = 2^{30}$, we have $n\delta(1) = 2^{-148}$, whereas $\delta(n) = 2^{-161}$. Computing exact $\delta(n)$ for Kyber and Saber would be more computationally exhaustive because there are many more possibilities for how the vector with the largest norm looks like when the coefficients are larger. But the secret key norm will still be tightly concentrated around its expected value and we will thus still have $\delta(n) < n\delta(1)$.

The n -user CPA advantage $\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}$. By a hybrid argument [6] one obtains the trivial bounds

$$\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \leq \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}} \leq n q_C \cdot \text{Adv}_{\text{PKE}}^{\text{IND-CPA}}.$$

Table 1: $\delta(n)$ for a basic LWE encryption scheme over \mathbb{Z}_q , for $q = 1409$, where the secret $\vec{s} \leftarrow \psi^{2048}$, where $\psi(0) = 1/2, \psi(\pm 1) = 1/4$, and the public key is $A \leftarrow \mathbb{Z}_q^{1024 \times 1024}, \vec{t} = [A \ I] \cdot \vec{s} \bmod q$. The ciphertext is constructed as $\vec{u} = [A^T \ I] \cdot \vec{r}^T, v = [\vec{t}^T \ 0^{1024}] \cdot \vec{r} + e + \lfloor q/2 \rfloor \mu$, where $\vec{r} \leftarrow \psi^{2048}, e \leftarrow \psi$.

n	1	2^{10}	2^{20}	2^{30}	2^{40}	2^{50}
$\log(\delta(n))$	-178	-170	-165	-161	-158	-155

For schemes based on prime-order groups (eg., ElGamal) we actually have $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \approx \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}$ by the well known random self reducibility of Diffie-Hellman tuples.

For lattice-based schemes, we do not have self-reductions that allow us to directly conclude that $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \approx \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}$ because we do not know of a way to create extra samples from the same distribution (i.e. either the MLWE one or uniform) even if these samples are allowed to have completely distinct secrets. Nevertheless, if we assume the hardness of MLWE as originally defined for the purpose of worst-case to average-case reductions [28, 30, 32] where the number of samples (using the same secret) is unlimited, then we can show that $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \approx \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}$. In particular, if we assume that distinguishing $\{(A_i, A_i \cdot \vec{s}_i + \vec{e}_i)\}_{1 \leq i \leq k}$ from uniform is hard for $k = \max(n, q_C)$,² then using the transformation from [2], one can argue that distinguishing the n public keys from uniform ones is as hard as MLWE. After replacing the public keys with random values, simulating the q_C queries to the n public keys can again be done with access to q_C samples of an MLWE problem and the transformation from [2]. And thus we have $\text{Adv}_{\text{PKE}}^{\text{IND-CPA}} \leq 2 \cdot \text{Adv}_{\text{PKE}}^{\text{MLWE}}$, which is also what we based the hardness of the IND-CPA scheme on. In short, we believe that in practice the MLWE problem with k samples is no easier than with 1 sample, and then the bound in our reduction is significantly tighter than the previously known one. And independently of the relationship between the number of MLWE samples, the new QROM bound for $\text{FO}_{\text{ID}(pk), m}^\perp$ is also noticeably better than the one for FO_m^\perp .

²Unless k is so large that the Arora-Ge [3] attack applies, it is not known that seeing more samples makes the MLWE problem any easier in practice. It should also be noted that even if such an attack existed, it would still not necessarily apply to increasing $\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}$. This is because the secrets in the $(n, q_C)\text{-IND-CPA}$ game are all distinct, and so it is a plausibly much harder problem than MLWE with many samples for the same secret. In particular, we do not know of any better algorithm for solving MLWE with many samples if the secrets are small and different for every sample.

Prefix collisions. We note that given a target public-key pk , it might be possible for an adversary to efficiently compute a different public key pk' and a matching sk' such that $\text{ID}(pk) = \text{ID}(pk')$. This “attack” does not violate any of our security claims, nor do we see any practical attack that could exploit this property.

But we stress that $\text{ID}(pk)$ should never be used as a fingerprint. Collision-prone public-key fingerprints are known to be insecure and have been successfully exploited in practice, for example in the context of PGP [29].

1.3 Impact on Efficiency

We now measure the effect of replacing the original FO transform used in the Kyber and Saber KEMs with our new transform $\text{FO}_{\text{ID}(pk),m}^{\perp}$ (see Figure 4). We implemented $\text{FO}_{\text{ID}(pk),m}^{\perp}$ on top of the underlying CPA-secure encryption schemes in the official AVX2-optimized implementations of Kyber and Saber, and performed benchmarks on an Intel Skylake CPU. Concretely, the numbers in Tables 2 and 3 are the medians of the cycle counts of 10000 executions of the key generation (**K**), encapsulation (**E**), and decapsulation (**D**) operations for either the original FO transform or the new one from this work.

For the 32-byte prefix $\text{ID}(pk)$ of the public key in Kyber and Saber one can take the seed ρ that is already of size 32 bytes and uniformly random in these schemes. Alternatively, for Kyber, the first 33 bytes of the bitpacked representation of the polynomial vector \vec{t} in the public key can also be used. This is sufficient since \vec{t} is given in the NTT basis and contains the independently random short error vector \vec{e} as an additive term. It follows from a simple Fourier analysis computation as in [4, 13] that the first few NTT coefficient of \vec{e} are close to uniform modulo $q = 3329$. Concretely, the first 22 coefficients have more than 256 bit of entropy and they occupy 33 bytes in the bitpacking that uses 12 bits per coefficient. Taking the prefix from \vec{t} instead of ρ has the advantage that this is still secure in the slightly modified but fully compatible variant of Kyber where users re-use ρ and hence the MLWE matrix **A**.

We observe that one obtains significant speed improvements throughout all parameter sets and variants. For example, the key generation and encapsulation of Kyber-512-90s are 56% and 66% faster, respectively, when using $\text{FO}_{\text{ID}(pk),m}^{\perp}$.³ There is no speed-up in decryption since in the original Kyber and Saber CCA transforms avoid the expensive full public key hash needed for the re-encryption during decapsulation by pre-computing this hash during key generation and storing it in the secret key. But there is still an expensive hashing operation in decapsulation that is responsible for a significant portion of the runtime. Namely, the hash for \tilde{K} in line 11 of the decapsulation algorithm in Figure 4, which includes the full ciphertext as input. The fake key \tilde{K} is only needed when re-encryption fails, so one could achieve a significant speed-up also in decapsulation by only computing \tilde{K} in the case of decryption failure. The resultant FO transform would effectively be a middle ground between implicit and explicit rejection as it would leak rejection only via a timing side-channel. So in an application where timing side-channels are absolutely of no concern, this

³It should be mentioned again that another reason that encapsulation has a larger increase is that we removed an additional hash of the ciphertext. This was already shown to be intuitively unnecessary in [10], and in this work we show that it is also unnecessary in the multi-user setting.

faster and fully compatible transformation can be used. It is well known that explicit rejection is secure in the ROM [21]. One can also use the techniques from this paper to show improved bounds in the multi-user/challenge setting. Moreover, Theorem 6.1 in [15] now gives a security proof for explicit rejection in the QROM, but unfortunately only with much worse security bounds. Given the recent progress in this area, it is natural to think that there will be better bounds in the future, decreasing the gap between implicit and explicit rejection. For these reasons we have also benchmarked this soft variant of explicit rejection. The resulting cycle counts are included in brackets in Tables 2 and 3 and we see that one would for example achieve a speed-up of 45% for the Kyber512-90s decapsulation runtime.

Our new transform also has a noticeably larger effect on Kyber than Saber because the CPA-secure scheme underlying Kyber is more efficient than its Saber counterpart. Thus the running time of Kyber with the original FO transform was much more dominated by hashing.

2 PRELIMINARIES

For $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$. For a set S , $|S|$ denotes the cardinality of S . For a finite set S , we denote the sampling of a uniform random element x by $x \xleftarrow{\$} S$. The min entropy of a discrete random variable X is defined as $H_{\infty}(X) = -\log(\max_x \Pr[X = x])$.

2.1 Cryptographic Definitions

PUBLIC-KEY ENCRYPTION. A public-key encryption scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ consists of three algorithms, and a finite message space \mathcal{M} . The key generation algorithm Gen outputs a key pair (pk, sk) , where pk also defines a finite randomness space $\mathcal{R} = \mathcal{R}(pk)$ as well as a ciphertext space \mathcal{C} . The encryption algorithm Enc , on input pk and a message $m \in \mathcal{M}$, outputs an encryption $c \xleftarrow{\$} \text{Enc}(pk, m)$ of m under the public key pk . If necessary, we make the used randomness of encryption explicit by writing $c := \text{Enc}(pk, m; r)$, where $r \in \mathcal{R}$. The decryption algorithm Dec , on input sk and a ciphertext c , outputs either a message $m = \text{Dec}(sk, c) \in \mathcal{M}$ or a special symbol $\perp \notin \mathcal{M}$ to indicate that c is not a valid ciphertext.

PKE has n -user correctness error $\delta(n)$ if

$$\mathbb{E} \left[\max_{j \in [n]} \max_{m \in \mathcal{M}} \Pr \left[\text{Dec}(sk_j, \text{Enc}(pk_j, m)) \neq m \right] \right] \leq \delta(n),$$

where the expectation is taken over $((pk_1, sk_1), \dots, (pk_n, sk_n)) \xleftarrow{\$} (\text{Gen})^n$. For $n = 1$ we obtain the single-user correctness definition $\delta := \delta(1)$ of [21]. Note that $\delta \leq \delta(n) \leq n\delta$. Furthermore, there exists schemes for which either $\delta(n) = \delta$ or $\delta(n) = n\delta$. PKE is weakly γ -spread [15] if

$$\mathbb{E} \left[\max_{m,c} \Pr \left[\text{Enc}_{pk}(m) = c \right] \right] \leq 2^{-\gamma},$$

where the expectation is taken over $(pk, sk) \xleftarrow{\$} \text{Gen}$.

We define n -user/ q_C -challenges IND-CPA (“passive”) security for PKE in terms of the advantage function of an adversary \mathcal{A} :

$$\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}(\mathcal{A}) := \left| \Pr[(n, q_C)\text{-IND-CPA}^{\mathcal{A}} \Rightarrow 1] - \frac{1}{2} \right|,$$

where game $(n, q_C)\text{-IND-CPA}$ is defined in Fig. 2.

Table 2: Median Skylake cycle counts of 10000 executions of Kyber and Saber using either the original CCA transform or the improved transform from this work. The “original” version of Kyber and Saber are the Round 3 submissions to the NIST post-quantum standardization process. The decapsulation cycle counts in brackets are for the “soft” explicit rejection variant that leaks rejection via a timing side-channel.

NIST Level		Kyber			Saber		
		Original	This Work	Speed-up	Original	This Work	Speed-up
1	K	23562	12883	45%	42169	36220	14%
	E	37144	16981	54%	57831	39232	32%
	D	28595	28529 (18332)	0% (36%)	57780	57806 (47516)	0% (18%)
3	K	40487	25272	38%	74577	64180	14%
	E	55726	27624	50%	95958	69304	28%
	D	43553	43442 (29660)	0% (32%)	95388	95301 (80847)	0% (15%)
5	K	55770	38815	30%	116178	102101	12%
	E	77011	40692	47%	142034	109203	23%
	D	61470	61473 (43194)	0% (30%)	142957	143090 (125589)	0% (12%)

Table 3: Median Skylake cycle counts of 10000 executions of the 90’s variants of Kyber and Saber using either the original CCA transform or the improved transform from this work. The “original” version of Kyber and Saber are the Round 3 submissions to the NIST post-quantum standardization process. The decapsulation cycle counts in brackets are for the “soft” explicit rejection variant that leaks rejection via a timing side-channel.

NIST Level		Kyber90s			uSaber90s		
		Original	This Work	Speed-up	Original	This Work	Speed-up
1	K	13994	6224	56%	24557	17294	30%
	E	23069	7894	66%	36544	22363	39%
	D	16917	16959 (9233)	0% (45%)	38156	38014 (30579)	0% (20%)
3	K	21783	10995	50%	37511	29881	20%
	E	33534	13137	61%	55436	36282	35%
	D	25014	24957 (14893)	0% (40%)	58395	58405 (47889)	0% (18%)
5	K	31576	18834	40%	59169	50796	14%
	E	46881	21404	54%	81187	57920	29%
	D	36190	36165 (22605)	0% (38%)	86142	86359 (72879)	0% (15%)

(n, q_C) -IND-CPA 01 for $j \in [n]$ 02 $(pk_j, sk_j) \xleftarrow{\$} \text{Gen}$ 03 $\vec{pk} \leftarrow (pk_1, \dots, pk_n)$ 04 $b \xleftarrow{\$} \{0, 1\}$ 05 $b' \xleftarrow{\$} \mathcal{A}^{\text{Chall}}(\vec{pk})$ 06 return $\llbracket b' = b \rrbracket$ $\text{Chall}(j, m_0, m_1)$ / max. q_C queries 07 return $\text{Enc}(pk_j, m_b)$
--

Figure 2: Game (n, q_C) -IND-CPA for PKE in the n -user/ q_C -challenges setting. Set \mathcal{L}_{M_j} describes the set of plaintext messages encrypted to user j .

KEY ENCAPSULATION MECHANISMS. A key encapsulation mechanism $\text{KEM} = (\text{Gen}, \text{Encaps}, \text{Decaps})$ consists of three algorithms and a finite key space \mathcal{K} similar to a PKE scheme, but Encaps does not take a message as input. The key generation algorithm Gen outputs a key pair (pk, sk) , where pk also defines a finite randomness space $\mathcal{R} = \mathcal{R}(pk)$ as well as a ciphertext space \mathcal{C} . The encapsulation algorithm Encaps takes as input a public-key pk and outputs a key encapsulation ciphertext c and a key k , that is $(c, k) \xleftarrow{\$} \text{Encaps}(pk)$. The decapsulation algorithm Decaps, on input sk and a ciphertext c , outputs either a key $k = \text{Decaps}(sk, c) \in \mathcal{K}$ or a special symbol $\perp \notin \mathcal{K}$ to indicate that c is not a valid ciphertext.

In terms of KEM’s security, we consider the n -user/ q_C -challenges IND-CCA advantage function of an adversary \mathcal{A} :

$$\text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}(\mathcal{A}) := \left| \Pr[(n, q_C)\text{-IND-CCA}_{\text{KEM}}^{\mathcal{A}} \Rightarrow 1] - \frac{1}{2} \right|$$

(n, q_C) -IND-CCA	Chall(j) / max. q_C queries
01 for $j \in [n]$	07 $(c, K_0) \xleftarrow{\$} \text{Encaps}(pk_j)$
02 $(pk_j, sk_j) \xleftarrow{\$} \text{Gen}$	08 $K_1 \xleftarrow{\$} \mathcal{K}$
03 $\vec{pk} \leftarrow (pk_1, \dots, pk_n)$	09 $\mathcal{L}_{C_j} := \mathcal{L}_{C_j} \cup \{c\}$
04 $b \xleftarrow{\$} \{0, 1\}$	10 return (c, K_b)
05 $b' \xleftarrow{\$} \mathcal{A}^{\text{Decaps, Chall}}(\vec{pk})$	Decaps($j, c \notin \mathcal{L}_{C_j}$)
06 return $\llbracket b = b' \rrbracket$	11 return Decaps(sk_j, c)

Figure 3: Game (n, q_C) -IND-CCA for KEM in the n -user/ q_C -challenges setting. Set \mathcal{L}_{C_j} describes the set of challenge ciphertexts for user j .

where game (n, q_C) -IND-CCA is defined in Fig. 3.

3 FUJISAKI-OKAMOTO TRANSFORMATION WITH PREFIX HASHING

Let $k, \lambda \in \mathbb{N}$ be integers. Let PKE be a public-key encryption scheme with message space \mathcal{M} , public-key space \mathcal{PK} , randomness space \mathcal{R} , and ciphertext space \mathcal{C} . We assume that the two spaces \mathcal{M} and $\{0, 1\}^k \times \mathcal{C}$ are disjoint. Let $\text{ID} : \mathcal{PK} \rightarrow \{0, 1\}^\ell$ be a fixed-output length function and let $F : \{0, 1\}^* \rightarrow \{0, 1\}^k \times \mathcal{R}$ be a hash function, where $F_1(X)$ is defined as the first k bits of $F(X)$. To PKE, ID, and F we associate the key encapsulation mechanism $\text{FO}_{\text{ID}(pk), m}^\ell[\text{PKE}, \text{ID}, F]$ as described in Fig. 4. Our transformation is essentially FO_m^ℓ of [21] with the difference that we feed $\text{ID}(pk)$ into the hash function F for domain separation. Note that with the identity function $\text{ID}(pk) = pk$, we recover $\text{FO}_{\text{ID}(pk), m}^\ell = \text{FO}_{pk, m}^\ell$; with $\text{ID}(pk) = \varepsilon$, we recover FO_m^ℓ .

Note that computing the pseudo-random key in line 11 of Decaps(sk, c) is only required in case c is inconsistent. We still recommend to always compute \tilde{K} because otherwise the system might be prone to a simple side-channel attack.

We now state the our main theorems about $\text{FO}_{\text{ID}(pk), m}^\ell[\text{PKE}, \text{ID}, F]$'s security in the ROM and QROM, respectively. In the concrete security statements we will use the following terms

- n -user correctness error $\delta(n)$
- (weak) γ -spreadness
- Min entropy ℓ of $\text{ID}(pk)$, i.e., $\ell := H_\infty(pk)$, where $(pk, sk) \xleftarrow{\$} \text{Gen}$
- Bit-length λ of the secret seed $s \in \{0, 1\}^\lambda$
- Max. number of (Q)ROM queries q_F
- Max. number of decapsulation queries q_D
- Max. number of challenge queries q_C

THEOREM 3.1 ((n, q_C) -IND-CPA OF PKE $\xrightarrow{\text{ROM}}$ (n, q_C) -IND-CCA OF KEM). *For any adversary \mathcal{A} against the (n, q_C) -IND-CCA security of KEM := $\text{FO}_{\text{ID}(pk), m}^\ell[\text{PKE}, \text{ID}, F]$ there exist adversaries \mathcal{B} and \mathcal{C} against (n, q_C) -IND-CPA of PKE (with roughly the same running*

time) such that $\text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}(\mathcal{A}) \leq$

$$2\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}(\mathcal{B}) + \frac{2(q_F + q_C)q_C}{|\mathcal{M}|} + \frac{q_F}{2^\lambda} + (q_F + q_D) \cdot \delta(n) + \frac{n^2}{2^\ell} \quad (1)$$

$$2\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}(\mathcal{C}) + \frac{2(q_F + q_C)q_C}{|\mathcal{M}|} + \frac{n(q_F + n)}{2^\lambda} + (q_F + q_D) \cdot n\delta + n \cdot q_D 2^{-\ell}. \quad (2)$$

Equation (1) of Theorem 3.1 will be proved in Section 4, equation (2) in Section 7. The proofs follow essentially the ones from [21], where we have to take extra care to allow for a reduction from multi-user/challenge IND-CPA security.

Note that bound (1) is meaningless for small values ℓ . Bound (2) is slightly weaker but independent of ℓ and therefore also holds for $\text{ID}(pk) = \varepsilon$.

THEOREM 3.2 ((n, q_C) -IND-CPA OF PKE $\xrightarrow{\text{QROM}}$ (n, q_C) -IND-CCA OF KEM). *For any quantum adversary \mathcal{A} against the (n, q_C) -IND-CCA security of KEM := $\text{FO}_{\text{ID}(pk), m}^\ell[\text{PKE}, \text{ID}, F]$ there exists a quantum adversary \mathcal{B} against (n, q_C) -IND-CPA of PKE (with roughly the same running time) such that $\text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}(\mathcal{A}) \leq$*

$$2\sqrt{q\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}(\mathcal{B})} + 4q\sqrt{\frac{q_C \cdot n}{|\mathcal{M}|} + \frac{n^2}{2^\ell} + 4(q_F + 1)\sqrt{\frac{n}{2^\lambda} + 16q^2\delta(n)} + \frac{q_C^2}{|\mathcal{M}|}},$$

where $q := q_F + q_D + 1$.

The proof of Theorem 3.2 is given in Section 5.

4 PROOF OF THEOREM 3.1

In this section we prove (1) of Theorem 3.1. The proof of (2) is similar and can be found in Section 7.

PROOF. Let \mathcal{A} be an adversary and consider the games given in Fig. 5.

GAME G_0 . This is the original (n, q_C) -IND-CCA game.

$$\left| \Pr[G_0^{\mathcal{A}} \Rightarrow 1] - \frac{1}{2} \right| = \text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}(\mathcal{A}).$$

GAME G_1 . In game G_1 an abort condition COLL is introduced, which aborts if there is a collision in the public-key identifiers $\text{ID}(pk_j)$. In case there no collision in the identifiers, we are able to identify each identifier $\text{ID}(pk_j)$ with a unique index j pointing to pk_j . This allows us to internally simulate

$$F(\text{ID}(pk_j), A) := \begin{cases} (H_j(m), G_j(m)) & A = m \\ K_j(s, c) & A = (s, c) \end{cases},$$

where G_j and H_j are internal random oracles, i.e., perfect random functions not accessible by \mathcal{A} . (This works because the two spaces \mathcal{M} and $\{0, 1\}^k \times \mathcal{C}$ are disjoint.) Since the two games are identical until COLL happens, we have by the birthday bound

$$\left| \Pr[G_0^{\mathcal{A}} \Rightarrow 1] - \Pr[G_1^{\mathcal{A}} \Rightarrow 1] \right| \leq \Pr[\text{COLL}] \leq \frac{n^2}{2^\ell},$$

where $l = H_\infty(\text{ID}(pk))$ is the min-entropy of $\text{ID}(pk)$ for $(pk, sk) \xleftarrow{\$} \text{Gen}$.

Gen'	Encaps(pk)	Decaps((sk, s), c)
01 $(pk, sk) \xleftarrow{\$} \text{Gen}$	05 $m \xleftarrow{\$} \mathcal{M}$	09 $m' \leftarrow \text{Dec}(sk, c)$
02 $s \xleftarrow{\$} \{0, 1\}^\lambda$	06 $(K, r) \leftarrow \text{F}(\text{ID}(pk), m)$	10 $(K, r) \leftarrow \text{F}(\text{ID}(pk), m')$
03 $sk' := (sk, s)$	07 $c \leftarrow \text{Enc}(pk, m; r)$	11 $\tilde{K} := \text{F}_1(\text{ID}(pk), s, c)$
04 return (pk, sk')	08 return (K, c)	12 if $m' = \perp$ or $\text{Enc}(pk, m'; r) \neq c$
		return \tilde{K}
		13 else return K

Figure 4: KEM = $\text{FO}_{\text{ID}(pk), m}^{\perp}[\text{PKE}, \text{ID}, \text{F}]$ with “implicit rejection” and “partial key hashing” built from PKE, F, and ID.

GAMES $G_0 - G_5$	$\text{F}(\text{id}, A)$	
01 for $j \in [n]$	31 if $\exists j \in [n] : \text{id} = \text{ID}(pk_j)$	/ $G_1 - G_5$
02 $(pk_j, sk_j) \xleftarrow{\$} \text{Gen}$	32 if $A \in \mathcal{M}$	/ $G_1 - G_5$
03 $s_j \xleftarrow{\$} \{0, 1\}^\lambda$	33 $m := A$	/ $G_1 - G_5$
04 $\vec{pk} \leftarrow (pk_1, \dots, pk_n)$	34 $(K, r) := (\text{H}_j(m), \text{G}_j(m))$	/ $G_1 - G_5$
05 $\exists i, j \in [n]$ s.t. $\text{ID}(pk_i) = \text{ID}(pk_j)$	35 if $A \in \{0, 1\}^\lambda \times C$	/ $G_1 - G_5$
06 COLL := true	36 $(s, c) := A$	/ $G_1 - G_5$
07 abort	37 $K := \text{K}_j(s, c)$	/ $G_1 - G_5$
08 $b \xleftarrow{\$} \{0, 1\}$	38 if K undefined: $K \xleftarrow{\$} \mathcal{K}$	
09 $b' \xleftarrow{\$} \mathcal{A}^{\text{Decaps}, \text{Chall}, \text{F}}(\vec{pk})$	39 if r undefined: $r \xleftarrow{\$} \mathcal{R}$	
10 return $\llbracket b' = b \rrbracket$	40 return (K, r)	
Chall (j)	$\text{G}_j(m)$	/ Internal random oracle
11 $m \xleftarrow{\$} \mathcal{M}$	41 $r \xleftarrow{\$} \mathcal{R}$	
12 $m \xleftarrow{\$} \mathcal{M} \setminus \mathcal{L}_{M_j}$	42 if $m \in \mathcal{L}_{M_j}$	/ G_5
13 $(K_0, r) \leftarrow \text{F}(\text{ID}(pk_j), m)$	43 QUERY := true	/ G_5
14 $\mathcal{L}_{M_j} := \mathcal{L}_{M_j} \cup \{m\}$	44 abort	/ G_5
15 $r \xleftarrow{\$} \mathcal{R}; K_0 \xleftarrow{\$} \{0, 1\}^k$	45 return r	
16 $K_1 \xleftarrow{\$} \{0, 1\}^k$	$\text{K}_j(s, c)$	/ Internal random oracle
17 $c \leftarrow \text{Enc}(pk_j, m; r)$	46 $K \xleftarrow{\$} \mathcal{K}$	
18 $\mathcal{L}_{C_j} := \mathcal{L}_{C_j} \cup \{c\}$	47 if $s = s_j$	/ $G_3 - G_5$
19 return (c, K_b)	48 BAD := true	/ $G_3 - G_5$
Decaps ($j, c \notin \mathcal{L}_{C_j}$)	49 abort	/ $G_3 - G_5$
20 $m' := \text{Dec}(sk_j, c)$	50 return K	
21 $(K, r) \leftarrow \text{F}(\text{ID}(pk_j), m')$	$\text{H}_j(m)$	/ Internal random oracle
22 if $m' = \perp$ or $\text{Enc}(pk_j, m'; r) \neq c$	51 $K \xleftarrow{\$} \mathcal{K}$	
23 $(K, r) \leftarrow \text{F}(\text{ID}(pk_j), s_j, c)$	52 if $m \in \mathcal{L}_{M_j}$	/ G_5
24 $K := \text{K}'_j(c)$	53 QUERY := true	/ G_5
25 return K	54 abort	/ G_5
Decaps ($j, c \in \mathcal{L}_{C_j}$)	55 $c' := \text{Enc}(pk_j, m; \text{G}_j(m))$	/ $G_4 - G_5$
26 if $\exists K$ s. th. $(c, K) \in \mathcal{L}_{D_j}$	56 if $\exists K'$ such that $(c', K') \in \mathcal{L}_{D_j}$	/ $G_4 - G_5$
27 return K	57 $K := K'$	/ $G_4 - G_5$
28 $K \xleftarrow{\$} \mathcal{K}$	58 else	/ $G_4 - G_5$
29 $\mathcal{L}_{D_j} := \mathcal{L}_{D_j} \cup \{(c, K)\}$	59 $\mathcal{L}_{D_j} := \mathcal{L}_{D_j} \cup \{(c', K)\}$	/ $G_4 - G_5$
30 return K	60 return K	

Figure 5: Games $G_0 - G_5$ for the proof of Theorem 3.1. The internal random oracles $\text{K}'_j, \text{H}_j, \text{G}_j, \text{K}_j$ are not accessible by the adversary. We assume wlog that F is only queried once on each value (id, A) .

GAME G_2 In game G_2 the challenge oracle Chall samples m from the set $\mathcal{M} \setminus \mathcal{L}_{M_j}$ instead of \mathcal{M} . This is necessary since $m_i = m_j$ implies $c_i = c_j$. Therefore we have $K_i = K_j$ in the “real world”, but the KEM keys in the “random world” are independent. Consequently, an adversary would be able distinguish between the “real” from “random” keys. Since there are at most q_C challenge queries we

have (by the birthday bound)

$$\left| \Pr \left[G_1^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[G_2^{\mathcal{A}} \Rightarrow 1 \right] \right| \leq \frac{q_C^2}{|\mathcal{M}|}.$$

GAME G_3 . In game G_3 we modify the Decaps oracle in lines 23 and 24 such that for an invalid ciphertext the key is defined as $\text{K}'_j(c)$, where K'_j is an independent internal random oracle. This remains unnoticed to adversary \mathcal{A} unless it queries $\text{K}_j(s_j, \cdot)$ for some $j \in [n]$. Since the seeds $s_j \in \{0, 1\}^\lambda$ are uniformly random

and information-theoretically hidden from the adversary, we have by the union bound

$$\left| \Pr[G_2^{\mathcal{A}} \Rightarrow 1] - \Pr[G_3^{\mathcal{A}} \Rightarrow 1] \right| \leq \frac{q_F}{2^\lambda}.$$

GAME G_4 . In game G_4 we simulate the decapsulation oracle $\text{Decaps}(j, \cdot)$ without knowledge of the secret key by patching the random oracles H_j in lines 55-59. Note that if PKE was perfectly correct, then the random oracle patching is also perfectly correct and therefore the two games would look identical in \mathcal{A} 's view.

The only bad case happens if G_j is queried on some m which induces a correctness error, that is $\text{Dec}(sk_j, \text{Enc}(pk_j; m; G_j(m))) \neq m$. More concretely, define the sets

$$\text{BAD}_j := \left\{ m \in \mathcal{M} \left| \begin{array}{l} m \neq m', \text{ where } c \leftarrow \text{Enc}(pk_j; m; G_j(m)); \\ m' \leftarrow \text{Dec}(sk_j, c) \end{array} \right. \right\}$$

Define the event CORR to be the event that \mathcal{A} makes an (implicit) query to $G_j(m)$ for some $m \in \text{BAD}_j$. Since there are at most $(q_F + q_D)$ explicit and implicit queries to G_j , we have

$$\Pr[\text{CORR}] \leq (q_F + q_D)\delta(n).$$

We now claim that

$$\left| \Pr[G_3^{\mathcal{A}} \Rightarrow 1] - \Pr[G_4^{\mathcal{A}} \Rightarrow 1] \right| \leq \Pr[\text{CORR}].$$

Let us analyze why G_3 and G_4 are identical conditioned on $\neg\text{CORR}$. Consider a query $\text{Decaps}(j, c)$ and define $m' := \text{Dec}(sk_j, c)$ and $c' := \text{Enc}(pk_j, m'; G_j(m'))$.

- **Case 1:** $m' = \perp$. H_j cannot be called on $m' = \perp$ and hence the KEM key of $K = \text{Decaps}(sk_j, c) = K'_j(c)$ in G_3 is identically distributed as $(c, K) \in \mathcal{L}_{D_j}$ in G_4 .
- **Case 2:** $m' \neq \perp \wedge c \neq c'$. Both games return a uniform random key K . The only way for \mathcal{A} to detect a difference between the two games is if it makes a query $H_j(m)$ such that $\text{Enc}(pk_j, m; G_j(m)) = c$. (In G_4 , $H_j(m)$ would return the same key K as in $\text{Decaps}(j, c)$, whereas in G_3 the two keys would be independent.) Since $c \neq c'$ we also have $m \neq m'$. But such a query $H_j(m)$ would internally involve the query $G_j(m)$ for $m \in \text{BAD}_j$.
- **Case 3:** $m' \neq \perp \wedge c = c'$. In game G_3 , $\text{Decaps}(j, c)$ returns $K = H_j(m')$, whereas G_4 first picks a uniform K and patches $H_j(m)$ to match K for all m that deterministically encrypt to the same c , i.e., all m satisfying $\text{Enc}(pk_j, m; G_j(m)) = c$. The only way for \mathcal{A} to detect a difference between the two games is to query H_j on some value $m \neq m'$ that also deterministically encrypts to the same c . But this also implies that $G_j(m)$ was queried for some $m \in \text{BAD}_j$.

GAME G_5 . In game G_5 we abort on queries of the form $H_j(m)$ or $G_j(m)$ for some challenge message $m \in \mathcal{L}_{M_j}$, in which case the event QUERY holds true. We have by the difference lemma,

$$\left| \Pr[G_4^{\mathcal{A}} \Rightarrow 1] - \Pr[G_5^{\mathcal{A}} \Rightarrow 1] \right| \leq \Pr[\text{QUERY}].$$

Note that in G_5 bit b is independent of the view of the adversary. We thus have

$$\Pr[G_5^{\mathcal{A}} \Rightarrow 1] = \frac{1}{2}.$$

We claim that

$$\Pr[\text{QUERY}] \leq 2 \cdot \left(\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}(\mathcal{B}) + \frac{q_F q_C}{|\mathcal{M}|} \right). \quad (3)$$

Summing up the inequalities yields the claimed bound, concluding the proof of the theorem.

We show (3) by giving an adversary \mathcal{B} against the (n, q_C) -IND-CPA security of PKE. Adversary \mathcal{B} from the (n, q_C) -IND-CPA challenger receives n public-keys and has access to a challenge encryption oracle. It runs \mathcal{B} on the public keys simulating the decapsulation oracle Decaps and the random oracle F as in G_4 and G_5 . On \mathcal{A} 's challenge query $\text{Chall}(j)$, \mathcal{B} proceeds as follows. It picks random $m_0 \xleftarrow{\$} \mathcal{M} \setminus \mathcal{L}_{M_{j,0}}$ and $m_1 \xleftarrow{\$} \mathcal{M} \setminus \mathcal{L}_{M_{j,1}}$ and adds m_0 to list $\mathcal{L}_{M_{j,0}}$ and m_1 to list $\mathcal{L}_{M_{j,1}}$. (If it is the first challenge query on input j it initializes the lists $\mathcal{L}_{M_{j,0}}$ and $\mathcal{L}_{M_{j,1}}$ to be empty.) Next, \mathcal{B} queries its own challenge oracle to obtain $c \leftarrow \text{Enc}(pk_j, m_b)$, where b is the (n, q_C) -IND-CPA's challenge bit. Finally, it returns (c, K) to \mathcal{A} , for a random key $K \in \{0, 1\}^k$. Note that this perfectly simulates the challenge oracle as in games G_4 and G_5 .

If, during the simulation of F , \mathcal{B} detects a query $H_j(m)$ or $G_j(m)$ for some $m \in \mathcal{L}_{M_{j,b'}}$, it returns b' and terminates. If no such query happens and \mathcal{A} terminates, \mathcal{B} returns a uniform bit b' .

Since all messages in the set $\mathcal{L}_{M_{j,1-b}}$ are information-theoretically hidden from \mathcal{A} , the probability that it queries $G_j(m)$ or $H_j(m)$ on some $m \in \mathcal{L}_{M_{j,1-b}}$ is bounded by $q_F q_C / |\mathcal{M}|$. Assume that this is not the case. We have

$$\begin{aligned} \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}(\mathcal{B}) + \frac{q_F q_C}{|\mathcal{M}|} &\geq \left| \Pr[b = b'] - \frac{1}{2} \right| \\ &= \left| \Pr[\text{QUERY}] + \frac{1}{2} \Pr[\neg\text{QUERY}] - \frac{1}{2} \right| \\ &= \frac{1}{2} \Pr[\text{QUERY}], \end{aligned}$$

which proves (3). \square

5 PROOF OF THEOREM 3.2

We refer to Section 6 for standard quantum notation. Our notation and the presentation of known results closely follows [23] and [22, Section 1.3].

We will now recall some QROM theorems that we will use during our proof of Theorem 3.2.

LEMMA 5.1 (GENERIC DISTINGUISHING PROBLEM WITH BOUNDED PROBABILITIES [23]). *Let X be a finite set, and let $\lambda \in [0, 1]$. For any (unbounded, quantum) algorithm \mathcal{A} issuing at most q quantum queries to F ,*

$$\left| \Pr[\text{GDPB}_{\lambda,0}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{GDPB}_{\lambda,1}^{\mathcal{A}} \Rightarrow 1] \right| \leq 8 \cdot \lambda \cdot (q+1)^2, \quad (4)$$

where games $\text{GDPB}_{\lambda,b}$ (for bit $b \in \{0, 1\}$) are defined in Fig. 6.

ONEWAY TO HIDING WITH SEMI-CLASSICAL ORACLES. In [1], Ambainis et al. defined semi-classical oracles that return a state that was measured with respect to one of the input registers. To any subset $S \subset X$, one can associate the following ‘‘semi-classical’’ oracle O_S^{SC} : Intuitively, O_S^{SC} collapses states taken from $\mathcal{H}_{X \times Y}$ to a state that contains only elements of either S or $X \setminus S$. To be more precise, O_S^{SC} takes as input a quantum state $|\psi\rangle, 0\rangle$ such that $|\psi\rangle \in \mathcal{H}_X \otimes \mathcal{H}_Y$.


```

GAME GDPB $_{\lambda, b}$ 
01  $(\lambda_x)_{x \in X} \leftarrow \mathcal{A}_1$ 
02 if  $\exists x \in X$  s.t.  $\lambda_x > \lambda$ 
03   return 0
04 if  $b = 1$ 
05   for all  $x \in X$ 
06      $O(x) \leftarrow B_{\lambda_x}$ 
07 else
08    $O := 0$ 
09  $b' \leftarrow \mathcal{A}_2^{(O)}$ 
10 return  $b'$ 

```

Figure 6: Generic distinguishing games GDPB $_{\lambda, b}$ with bounded maximal Bernoulli parameter $\lambda \in [0, 1]$.

O_S^{SC} first measures the X -register with respect to the projectors $M_1 := \sum_{x \in S} |x\rangle\langle x|$ and $M_0 := \sum_{x \notin S} |x\rangle\langle x|$, and then initialises the last register to $|b\rangle$ for the measured bit b . Consequently, $|\psi, 0\rangle$ collapses to either a state $|\psi', 1\rangle$ such that the X -register of $|\psi'\rangle$ only contains elements of S , or a state $|\psi', 0\rangle$ such that the X -register of $|\psi'\rangle$ only contains elements of $X \setminus S$.

To a quantum-accessible oracle O and a subset $S \subset X$, one can furthermore associate oracle $O \setminus S$ which first queries O_S^{SC} and then O . Let FIND denote the event that O_S^{SC} ever returns a state $|\psi', 1\rangle$. Unless FIND occurs, the outcome of $O \setminus S$ is independent of the values $O(x)$ for all $x \in S$, which is why $O \setminus S$ is also called a ‘‘punctured’’ oracle.

We will now restate several ‘‘semi-classical one-way to hiding’’ theorems from [1]. While [1] consider adversaries that might execute parallel oracle invocations, and therefore differentiate between query depth d and number of queries q , we use the upper bound $q \geq d$ for the sake of simplicity. Lemma 5.2 is a simplification of [1, Thm. 1: ‘‘Semi-classical O2H’’], and Eq. (6) (Eq. (7)) of Lemma 5.3 is a simplification of [1, Thm. 2: ‘‘Search in semi-classical oracle’’] ([1, Cor. 1]).

LEMMA 5.2. *Let $S \subset X$ be random. Let $O_1, O_2 \in Y^X$ be random functions such that $O_1(x) = O_2(x)$ for all $x \in X \setminus S$, and let z be a random bitstring. (S, O_1, O_2 , and z may have an arbitrary joint distribution.) For $i \in \{1, 2\}$, let*

$$p_i := \Pr[1 \leftarrow A^{(O_i)}(z)],$$

and let

$$p_{\text{FIND}} := \Pr[b \leftarrow A^{(O \setminus S)}(z) : \text{FIND}].$$

For all quantum algorithms A with binary output, issuing at most q queries, we have that

$$|p_1 - p_2| \leq 2 \cdot \sqrt{(q+1) \cdot p_{\text{FIND}}}. \quad (5)$$

LEMMA 5.3. *Let $S \subset X$ be random, let O be a random function, and let z be a random bitstring. (S, O , and z may have an arbitrary joint distribution.) Let*

$$p_{\text{FIND}} := \Pr[b \leftarrow A^{(O \setminus S)}(z) : \text{FIND}].$$

Then, for all quantum algorithms A with binary output issuing at most q queries, we have that

$$p_{\text{FIND}} \leq 4q \cdot \Pr[x \leftarrow B(z) : x \in S], \quad (6)$$

where B is the algorithm that, on input z , chooses $i \stackrel{\$}{\leftarrow} \{1, \dots, q\}$, runs $A^{(O)}$ until (just before) the i -th query, measures its query input register in the computational basis and outputs the measurement outcome.

If $S := \{x_1, \dots, x_n\}$ for $x_1, \dots, x_n \stackrel{\$}{\leftarrow} X$, and S and z are independent, we have that

$$p_{\text{FIND}} \leq \frac{4q|S|}{|X|}. \quad (7)$$

Furthermore, if $S := \{(y_1, x_1), \dots, (y_n, x_n)\}$ where $x_i \stackrel{\$}{\leftarrow} X \setminus \{x_1, \dots, x_{i-1}\}$ and $\{y_1, \dots, y_n\} \subset Y$ independent of $U := \{x_1, \dots, x_n\}$ we have for U independent of z

$$p_{\text{FIND}} \leq \frac{4q|S|}{|X|}. \quad (8)$$

We will now prove an additional helper lemma.

LEMMA 5.4 (QROM MULTI-USER PRF). *For $i \in [n]$, let $p_i \in \{0, 1\}^Y$ arbitrarily subject to $p_i \neq p_j$, when $i \neq j$. Define $\vec{p} := (p_1, \dots, p_n)$. Let H, H_1, \dots, H_n be independent random oracles with $H: \{0, 1\}^Y \times \{0, 1\}^\lambda \times X \rightarrow \mathcal{Y}'$ and $H_i: X \rightarrow \mathcal{Y}'$, then for all quantum algorithms \mathcal{A} issuing at most q quantum queries to H and arbitrarily many queries to H_i with $i \in [n]$, we have*

$$\left| \Pr \left[\mathcal{A}^{(H, H(p_1, s_1, \cdot), \dots, H(p_n, s_n, \cdot))}(\vec{p}) = 1 \right] - \Pr \left[\mathcal{A}^{(H, H_1, \dots, H_n)}(\vec{p}) = 1 \right] \right| \leq 4(q+1) \sqrt{\frac{n}{2^\lambda}},$$

where the probabilities are taken over H, H_1, \dots, H_n and $s_1, \dots, s_n \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$, and the internal randomness of \mathcal{A} .

PROOF. The proof follows from combining Lemma 5.2 with Lemma 5.3 and the overall proof idea is very similar to the one of [10, Corollary 1]. The adversary’s goal is to distinguish quantum access to H and additional classical access to $(H(p_1, s_1, \cdot), \dots, H(p_n, s_n, \cdot))$ from quantum access to H and additional classical access to a collection (H_1, \dots, H_n) of independent random oracles. This is the same as distinguishing $(H, H^{(p_1, s_1, \cdot)} \rightarrow H_1(\cdot), \dots, H^{(p_n, s_n, \cdot)} \rightarrow H_n(\cdot))$ from (H, H_1, \dots, H_n) , and the set of differences between these oracles is $S := \{(p_j, s_j) \mid j \in [n]\} \times X$. According to Lemma 5.2, the distinguishing advantage is at most $2\sqrt{(q+1) \cdot p_{\text{FIND}}}$, and we can apply Eq. (8) from Lemma 5.3 to obtain that $p_{\text{FIND}} \leq 4q \cdot \frac{n}{2^\lambda}$. It remains to note that for $i \neq j$ we have $(p_i, s_i) \neq (p_j, s_j)$. \square

5.1 Proof

With these results at hand, we can finally proceed to the proof of Theorem 3.2. Its overall structure is very similar to recent QROM proofs of IND-CCA security [10, 23–27, 33]: we first change the game such that the decapsulation oracle can be simulated without knowledge of the collection of secret keys, which is achieved as usual by plugging encryption into the random oracle (in our case for multiple users), and then apply one-way to hiding to argue key indistinguishability.

PROOF OF THEOREM 3.2. Let \mathcal{A} be an adversary and consider the games given in Fig. 7 and Fig. 9.

GAMES $G_0 - G_7$		$F(\text{id}, A)$	
01 for $j \in [n]$		25 if $A \in \mathcal{M}$	
02 $(pk_j, sk_j) \xleftarrow{\$} \text{Gen}$		26 $m := A$	
03 $s_j \xleftarrow{\$} \{0, 1\}^\lambda$		27 return $(H(\text{id}, m), G(\text{id}, m))$	
04 $\vec{pk} \leftarrow (pk_1, \dots, pk_n)$		28 else if $A \in \{0, 1\}^\lambda \times C$	
05 $\exists i, j \in [n]$ s.t. $\text{ID}(pk_i) = \text{ID}(pk_j)$	/ G_1 - G_7	29 $(s, c) := A$	
06 COLL := true	/ G_1 - G_7	30 return $(K(\text{id}, s, c), L(\text{id}, s, c))$	
07 abort	/ G_1 - G_7	31 return $L'(\text{id}, A)$	
08 $b \xleftarrow{\$} \{0, 1\}$		$H(\text{id}, m)$	/ Internal Random Oracle
09 $b' \xleftarrow{\$} \mathcal{A}^{\text{Decaps, Chall, F }}(\vec{pk})$		32 if $\exists j \in [n] : \text{id} = \text{ID}(pk_j)$	/ G_1 - G_7
10 return $\llbracket b' = b \rrbracket$		33 return $H_j(m)$	/ G_1 - G_7
		34 return $H_0(\text{id}, m)$	
Chall (j)		$H_j(m)$	/ Internal Random Oracle
11 $m \xleftarrow{\$} \mathcal{M}$		35 return $H_j^1(m)$	/ G_0 - G_4
12 $m \xleftarrow{\$} \mathcal{M} \setminus \mathcal{L}_{M_j}$	/ G_2 - G_7	36 return $H_j^2(\text{Enc}(pk_j, m; G'_j(m)))$	/ G_5 - G_6
13 $\mathcal{L}_{M_j} := \mathcal{L}_{M_j} \cup \{m\}$		37 return $H_j^2(\text{Enc}(pk_j, m; G_j(m)))$	/ G_7
14 $(K_0, r) \leftarrow F(\text{ID}(pk_j), m)$		$G(\text{id}, m)$	/ Internal Random Oracle
15 $K_1 \xleftarrow{\$} \{0, 1\}^k$		38 if $\exists j \in [n] : \text{id} = \text{ID}(pk_j)$	/ G_1 - G_7
16 $c \leftarrow \text{Enc}(pk_j, m; r)$		39 return $G_j(m)$	/ G_1 - G_3, G_7
17 $\mathcal{L}_{C_j} := \mathcal{L}_{C_j} \cup \{c\}$		40 return $G'_j(m)$	/ G_4 - G_6
18 return (c, K_b)		41 return $G_0(\text{id}, m)$	
Decaps ($j, c \notin \mathcal{L}_{C_j}$)	// G_0 - G_5	$G'_j(m)$	/ Internal Random Oracle
19 $m' := \text{Dec}(sk_j, c)$		42 return $\text{Sample}(\mathcal{R} \setminus \mathcal{R}_{\text{bad}}(pk_j, sk_j, m); R(j, m))$	
20 $(K, r) \leftarrow F(\text{ID}(pk_j), m')$		Decaps ($j, c \notin \mathcal{L}_{C_j}$)	// G_6 - G_7
21 if $m' = \perp$ or $\text{Enc}(pk_j, m'; r) \neq c$		43 return $H_j^2(c)$	
22 $(K, r) \leftarrow F(\text{ID}(pk_j), s_j, c)$			
23 $K := K'_j(c)$	/ G_3 - G_5		
24 return K			

Figure 7: Games $G_0 - G_7$ for the proof of Theorem 3.2. The internal random oracles $H, H_0, G, G_0, G_j, H_j^1, H_j^2, K$ and K'_j, L', R are not directly accessible to the adversary. The notation $|F\rangle$ denotes that the adversary \mathcal{A} has quantum access to the random oracle F . The output range of F is $\{0, 1\}^k \times \mathcal{R}$.

GAME G_0 . This is the original (n, q_C) -IND-CCA game. We have

$$\left| \Pr \left[G_0^{\mathcal{A}} \Rightarrow 1 \right] - \frac{1}{2} \right| = \text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}(\mathcal{A}).$$

GAME G_1 . In game G_1 we abort on a public-key identifier collision (line 05-line 07). Like in the ROM proof, we can now identify every public-key identifier $\text{ID}(pk_j)$ with its index j . This allows us to simulate the random oracles, when called on such an identifier and $m := A \in \mathcal{M}$, via

$$F(\text{ID}(pk_j), m) := (H_j(m), G_j(m)), \quad (9)$$

where H_j and G_j are independent internal random oracles. We stress the importance of this step. First, it will allow us in later steps of the proof to replace G_j with a random oracle G'_j which only samples “good randomness” with respect to the public-key pk_j . Second, it is essential in order to simulate the $\text{Decaps}(j, \cdot)$ oracle without knowledge of the secret-key. Since the identification of Eq. 9 is only a conceptual change, by the birthday bound and the ℓ -bit min-entropy of $\text{ID}(pk)$ we have

$$\left| \Pr \left[G_0^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[G_1^{\mathcal{A}} \Rightarrow 1 \right] \right| \leq \Pr[\text{COLL}] \leq \frac{n^2}{2^\ell}.$$

GAME G_2 In game G_2 the challenge oracle samples m from $\mathcal{M} \setminus \mathcal{L}_{M_j}$ instead of \mathcal{M} , since on a message collision an adversary can easily differentiate between the real or random world. By the birthday bound we have

$$\left| \Pr \left[G_1^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[G_2^{\mathcal{A}} \Rightarrow 1 \right] \right| \leq \frac{q_C^2}{|\mathcal{M}|}.$$

GAME G_3 . In game G_3 we substitute $F(\text{ID}(pk_j), s_j, c)$ in line 22, whose first k bits evaluate to $K(\text{ID}(pk_j), s_j, c)$ in line 30, with $K'_j(c)$ in line 23, where K'_j is an internal independent random oracle for every $j \in [n]$. With a straightforward reduction, we can apply Lemma 5.4. Concretely, we identify H with K , H_j with K'_j , p_j with $\text{ID}(pk_j)$ and the uniform secrets $s_j \in \{0, 1\}^\lambda$ from Lemma 5.4 with the scheme’s secret seeds s_j used in the game. We obtain

$$\left| \Pr \left[G_2^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[G_3^{\mathcal{A}} \Rightarrow 1 \right] \right| \leq 4(q_F + 1) \sqrt{\frac{n}{2^\lambda}}.$$

GAME G_4 . In game G_4 we switch the random oracle G to an oracle which samples only “good” randomness, meaning that no decryption failure of proper encryptions can possibly occur anymore. To

make this more formal, for fixed $(pk, sk) \in \text{supp}(\text{Gen})$ and $m \in \mathcal{M}$, let

$$\mathcal{R}_{\text{bad}}(pk, sk, m) := \{r \in \mathcal{R} : \text{Dec}(sk, \text{Enc}(pk, m; r)) \neq m\}$$

in order to define $\delta(pk, sk, m) := |\mathcal{R}_{\text{bad}}(pk, sk, m)| / |\mathcal{R}|$. The modified oracle G can now be defined as follows: We will still let G coincide with G_0 anywhere but on $\{\text{ID}(pk_j) \mid j \in [n]\} \times \mathcal{M}$. Whereas for any index j , however, $G(\text{ID}(pk_j), \cdot)$ was defined until this game by as an oracle G_j which has range \mathcal{R} (line 39), we now replace each G_j with a random oracle G'_j that has range $\mathcal{R} \setminus \mathcal{R}_{\text{bad}}(pk_j, sk_j, m)$ instead in line 40. We will now argue that distinguishing game G_3 from G_4 can be reduced to distinguishing the GDPB games (see Lemma 5.1). Consider the quantum distinguisher $\mathcal{D} = (\mathcal{D}_1, \mathcal{D}_2)$ given in Fig. 8. After choosing its Bernoulli parameters $\lambda(j, m) := \delta(pk_j, sk_j, m)$, \mathcal{D} is provided with access to oracle O which is either the constant 0 function (if \mathcal{D} is run in game $\text{GDPB}_{\lambda, 0}$) or distributed according to the chosen Bernoulli parameters (in game $\text{GDPB}_{\lambda, 1}$). Note that $\delta(n)$ serves as an upper bound for these parameters. \mathcal{D} perfectly simulates game G_3 if run in game $\text{GDPB}_{\lambda, 1}$, and game G_4 if run in $\text{GDPB}_{\lambda, 0}$. Identifying λ with $\delta(n)$ and q with $q_F + q_D$, we can apply Lemma 5.1 to obtain

$$\left| \Pr[G_3^{\mathcal{A}} \Rightarrow 1] - \Pr[G_4^{\mathcal{A}} \Rightarrow 1] \right| \leq 8(q_F + q_D + 1)^2 \delta(n).$$

GAME G_5 . In game G_5 , we modify the random oracle $H_j(m)$ such that it returns $H_j^2(\text{Enc}(pk_j, m; G'_j(m)))$ in line 36 instead of $H_j^1(m)$ in line 35. Here, H_j^1 and H_j^2 are independent internal random oracles. Since we use the modified G that only samples “good” randomness, the mapping $\text{Enc}(pk_j, \cdot; G'_j(\cdot))$ is injective for all $j \in [n]$. Furthermore, since $\text{ID}(pk_i) \neq \text{ID}(pk_j)$ for all $i \neq j$, we have that the following function h is injective:

$$h(\text{id}, m) := \begin{cases} (\text{id}, \text{Enc}(pk_j, m; G'_j(m))) & \text{if } \exists j \in [n] \text{ s.t. } \text{id} = \text{ID}(pk_j) \\ (\text{id}, m) & \text{else} \end{cases}. \quad (10)$$

We can now rewrite H in Game G_5 as $H = H^2 \circ h$, where we define the internal random oracle H^2 as

$$H^2(\text{id}, m) := \begin{cases} H_j^2(m) & \text{if } \exists j \in [n] \text{ s.t. } \text{id} = \text{ID}(pk_j) \\ H_0(\text{id}, m) & \text{else} \end{cases}.$$

After these changes, H equals $H^2 \circ h$ for an internal random oracle H^2 and an injective function h . Therefore, H is still a random oracle in G_5 and hence perfectly indistinguishable from the one of G_4 . Thus, we have shown

$$\Pr[G_4^{\mathcal{A}} \Rightarrow 1] = \Pr[G_5^{\mathcal{A}} \Rightarrow 1].$$

GAME G_6 . In game G_6 , we start to simulate the decapsulation oracle without knowledge of the secret keys. We modify $\text{Decaps}(j, c)$ to return $H_j^2(c)$ in line 43. Let $m' := \text{Dec}(sk_j, c)$ and consider the following two cases:

- Case 1: The ciphertext c lies in the image of encryption, that is $\text{Enc}(pk_j, m'; G'_j(m')) = c$. In this case we have

$$\begin{aligned} H_j^2(c) &= H_j^2(\text{Enc}(pk_j, m'; G'_j(m'))) \\ &= (H^2 \circ h)(\text{ID}(pk_j), m') \\ &= H(\text{ID}(pk_j), m') \\ &= \text{Decaps}(j, c). \end{aligned}$$

Therefore the oracles in G_5 and G_6 are identical.

- Case 2: The ciphertext c lies not in the image of encryption, that is $\text{Enc}(pk_j, m'; G'_j(m')) \neq c$. Here the adversary gets to see $K'_j(c)$ in G_5 and $H_j^2(c)$ in G_6 for internal random oracles K'_j and H_j^2 . In both games the adversary gets access to H^2 indirectly through $H(\text{ID}(pk_j), m)$ forcing m through the encryption algorithm. Therefore, the adversary can not access H^2 for ciphertexts lying outside the image space of encryption. Therefore no adversary can distinguish between $K'_j(c)$ in G_5 and $H_j^2(c)$.

We just have shown

$$\Pr[G_5^{\mathcal{A}} \Rightarrow 1] = \Pr[G_6^{\mathcal{A}} \Rightarrow 1].$$

GAME G_7 . In game G_7 , we switch back to using the original random oracle G in lines 39 and 37. With essentially the same argument as in the game-hop from G_3 to G_4 , we have

$$\left| \Pr[G_6^{\mathcal{A}} \Rightarrow 1] - \Pr[G_7^{\mathcal{A}} \Rightarrow 1] \right| \leq 8(q_F + q_D + 1)^2 \delta(n).$$

What we have achieved so far is that the decapsulation oracle is simulatable by a reduction that does not know the secret keys, we can now apply the one-way to hiding theorems in order to reduce key indistinguishability to an attacker against the underlying encryption scheme.

GAME G_8 . In game G_8 we make the following conceptual change, in order to be able to apply the one-way to hiding theorems: Instead of sampling the challenge plaintexts and computing the corresponding challenge keys $K_{j,i}$ on the fly in the Chall oracle (like in G_7), we now precompute these values before running the adversary, and use them in the Chall oracle. Since this change is purely conceptual, we have that

$$\Pr[G_7^{\mathcal{A}} \Rightarrow 1] = \Pr[G_8^{\mathcal{A}} \Rightarrow 1].$$

GAME G_9 . In game G_9 , we decouple the oracle values on the challenge plaintexts from the values that are used to compute the challenge ciphertexts and the challenge keys, see Fig. 9. In game G_9 , the adversary’s view is independent of b and

$$\Pr[G_9^{\mathcal{A}} \Rightarrow 1] = \frac{1}{2}.$$

We have reprogrammed F on the set $\{(\text{ID}(pk_j), m_{j,i}) \mid j \in [n], i \in [q_C]\}$, and we can now first apply Lemma 5.2 in order to argue that

$$\left| \Pr[G_8^{\mathcal{A}} \Rightarrow 1] - \Pr[G_9^{\mathcal{A}} \Rightarrow 1] \right| \leq 2\sqrt{(q_F + q_D + 1) \cdot \Pr[\text{FIND}]}.$$

Here $\Pr[\text{FIND}]$ denotes the probability that when game G_9 is run with the punctured counterpart $F \setminus S$ of F instead of F , $F \setminus S$ ever collapses the input register of one of the random oracle queries to

$\underline{\mathcal{D}_1}$ 01 for $j \in [n]$ 02 $(pk_j, sk_j) \stackrel{\$}{\leftarrow} \text{Gen}$ 03 for $m \in \mathcal{M}$ 04 $\lambda(j, m) := \delta(pk_j, sk_j, m)$ 05 return $((\lambda(j, m)))_{j \in [n], m \in \mathcal{M}}$ $\underline{G_j(m)}$ 06 if $O(j, m) = 0$ 07 $G_j(m) := \text{Sample}(\mathcal{R} \setminus \mathcal{R}_{bad}(pk_j, sk_j, m); R(j, m))$ 08 else 09 $G_j(m) := \text{Sample}(\mathcal{R}_{bad}(pk_j, sk_j, m); R(j, m))$ 10 return $G_j(m)$	$\underline{\mathcal{D}_2^{(0)}}$ 11 for $j \in [n]$ 12 $(pk_j, sk_j) \stackrel{\$}{\leftarrow} \text{Gen}$ 13 $s_j \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ 14 $\vec{pk} \leftarrow (pk_1, \dots, pk_n)$ 15 $\exists i, j \in [n]$ s.t. $\text{ID}(pk_i) = \text{ID}(pk_j)$: abort 16 $b \stackrel{\$}{\leftarrow} \{0, 1\}$ 17 $b' \stackrel{\$}{\leftarrow} \mathcal{A}^{\text{Decaps, Chall, F}}(\vec{pk})$ 18 return $\llbracket b' = b \rrbracket$ $\underline{G(\text{id}, m)}$ 19 if $\exists j \in [n] : \text{id} = \text{ID}(pk_j)$ 20 return $G_j(m)$ 21 return $G_0(\text{id}, m)$
--	--

Figure 8: Distinguisher $\mathcal{D} = (\mathcal{D}_1, \mathcal{D}_2)$ for the generic distinguishing problem with bounded probabilities. The internal random oracle R is used to provide the random coins of the sampling procedure. The oracles Decaps , Chall and F are simulated as in G_3 (or equivalently G_4) of Fig. 7.

$\underline{G_7-G_{10}}$ 01 for $j \in [n]$ 02 $(pk_j, sk_j) \stackrel{\$}{\leftarrow} \text{Gen}$ 03 $s_j \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ 04 $\vec{pk} \leftarrow (pk_1, \dots, pk_n)$ 05 $\exists i, j \in [n]$ s.t. $\text{ID}(pk_i) = \text{ID}(pk_j)$ 06 COLL := true 07 abort 08 for $j \in [n]$ 09 for $i \in [q_C]$ 10 $m_{j,i} \stackrel{\$}{\leftarrow} \mathcal{M} \setminus \mathcal{L}_{M_j}$ 11 $\mathcal{L}_{M_j} := \mathcal{L}_{M_j} \cup \{m_{j,i}\}$ 12 $(K_{j,i}, r_{j,i}) \leftarrow F(\text{ID}(pk_j), m_{j,i})$ 13 $F(\text{ID}(pk_j), m_{j,i}) \stackrel{\$}{\leftarrow} \{0, 1\}^k \times \mathcal{R}$ 14 $b \stackrel{\$}{\leftarrow} \{0, 1\}$ 15 $b' \stackrel{\$}{\leftarrow} \mathcal{A}^{\text{Decaps, Chall, F}}(\vec{pk})$ 16 return $\llbracket b' = b \rrbracket$	$\underline{\text{Chall}(j)} \quad / G_7$ 17 $m \stackrel{\$}{\leftarrow} \mathcal{M} \setminus \mathcal{L}_{M_j}$ 18 $\mathcal{L}_{M_j} := \mathcal{L}_{M_j} \cup \{m\}$ 19 $(K_0, r) \leftarrow F(\text{ID}(pk_j), m)$ 20 $c \leftarrow \text{Enc}(pk_j, m; r)$ 21 $K_1 \stackrel{\$}{\leftarrow} \{0, 1\}^k$ 22 $\mathcal{L}_{C_j} := \mathcal{L}_{C_j} \cup \{c\}$ 23 return (c, K_b) $\underline{\text{Chall}(j)} \quad / G_8-G_{10}$ 24 $K_0 := K_{j, \text{ctr}_j}$ 25 $c \leftarrow \text{Enc}(pk_j, m_{j, \text{ctr}_j}; r_{j, \text{ctr}_j})$ 26 $c \stackrel{\$}{\leftarrow} \text{Enc}(pk_j, 0)$ 27 $\mathcal{L}_{C_j} := \mathcal{L}_{C_j} \cup \{c\}$ 28 $\text{ctr}_j := \text{ctr}_j + 1$ 29 $K_1 \stackrel{\$}{\leftarrow} \{0, 1\}^k$ 30 return (c, K_b) $\quad \quad \quad / G_8-G_9$ $\quad \quad \quad / G_{10}$
---	--

Figure 9: Games $G_7 - G_{10}$. Apart from the reprogramming in line 13, all oracles except of the Chall oracle remain defined as in game G_7 (see Fig. 7). For $j \in [n]$, we initialize the counters ctr_j to 1.

an element of S . In order to upper bound $\Pr[\text{FIND}]$, we dissect S into the set S_1 of (identifiers and) messages that were actually used by Chall , and the set S_2 of (identifiers and) messages that were not used by Chall . More formally, we define the set

$$S_1 := \{(\text{ID}(pk_j), m_{j,i}) \mid j \in [n], i \in [q_{C,j}]\},$$

where $q_{C,j}$ denotes the number of queries that were issued to Chall on index j , and we let $S_2 := S \setminus S_1$. Accordingly, we can now also dissect event FIND into the event FIND_1 that a query was collapsed to an element of S_1 , and into the event FIND_2 that a query was collapsed to an element of S_2 . We have that

$$\Pr[\text{FIND}] \leq \Pr[\text{FIND}_1] + \Pr[\text{FIND}_2].$$

Note that the suffixes of the elements in S_2 are uniformly random (without repetition), and that these elements were used nowhere in the game except for the puncturing of F . In particular, they are independent of \mathcal{A} 's input. Since S_2 consists of $q_C \cdot (n-1)$ many elements, we can now apply Eq. (8) of Lemma 5.3 to argue that

$$\Pr[\text{FIND}_2] \leq \frac{4(q_F + q_D) \cdot q_C \cdot (n-1)}{|\mathcal{M}|}.$$

It remains to upper bound $\Pr[\text{FIND}_1]$, for which we will first give a reduction to (n, q_C) -IND-CPA. A (n, q_C) -IND-CPA reduction \mathcal{B} can sample the challenge plaintexts and the plaintexts from S_2 before running the adversary, it can hence perfectly simulate the puncturing. Whenever our reduction \mathcal{B} receives a challenge query from \mathcal{A} on an index j , it queries its own challenge oracle on $j, m_0 :=$

$m_{j, \text{ctr}_j}, m_1 := 0$. It uses the response of CHAL to simulate either game G_9 with punctured oracles (if the (n, q_C) -IND-CPA bit is 0) or game G_{10} with punctured oracles (if the (n, q_C) -IND-CPA bit is 1), we hence have

$$\Pr[\text{FIND}_1] \leq \text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}(\mathcal{B}) + \Pr[\text{FIND}_1 \text{ in } G_{10} \text{ with punctured oracle}] .$$

Since in game G_{10} , the suffixes of the elements in the set S_1 are also independent of A 's input, and since S_1 has q_C many elements, we can now apply Eq. (8) of Lemma 5.3 again to argue that

$$\Pr[\text{FIND}_1 \text{ in } G_{10} \text{ with punctured oracle}] \leq \frac{4(q_F + q_D) \cdot q_C}{|\mathcal{M}|} .$$

Adding the bounds yields the claimed bound in the theorem, concluding our proof.

REMARK ON THE RUNNING TIME OF THE REDUCTIONS. The running time of the reductions depends on the implementation of data structures. Naively, the running time is about that of the CCA adversary, plus additional quadratic terms in the number of hashing, challenge and decapsulation queries. Using techniques from [19] one can generate for example the plaintexts by means of a (strong) pseudorandom permutation/ideal cipher (e.g., $m_i = \text{PRP}(i)$). Then the reduction only needs to keep track of a counter, which it can use to implement set insertions and set membership tests efficiently. In that case (ignoring logarithmic terms), the running time of the reductions becomes about that of the CCA adversary (assuming that evaluation and inversion of the PRP and other primitives as the encryption algorithm is in $O(1)$). Similar statements hold for the quantum case, except that a fully-quantum secure PRP or quantum-accessible ideal cipher is needed, since the adversary gets quantum access through the quantum random oracle. Alternatively, using balanced search trees, one should get similar running times as using the PRP (up to logarithmic factors), but larger memory usage. Also note that \mathcal{B} has slightly worse running times (factor n on queries to F) than \mathcal{A} in Theorem 3.1, which is an argument for public-key/prefix hashing. \square

ACKNOWLEDGMENTS

We would like to thank Mike Hamburg for his helpful comments. Julien Duman was supported by the EU H2020 PROMETHEUS project 780701. Eike Kiltz was supported by the BMBF iBlockchain project, the EU H2020 PROMETHEUS project 780701, DFG SPP 1736 Big Data, and the DFG Cluster of Excellence 2092 CASA. Vadim Lyubashevsky and Gregor Seiler are supported by the EU H2020 ERC Project 101002845 "PLAZA".

REFERENCES

- [1] Andris Ambainis, Mike Hamburg, and Dominique Unruh. 2019. Quantum Security Proofs Using Semi-classical Oracles. In *CRYPTO 2019, Part II (LNCS)*, Alexandra Boldyreva and Daniele Micciancio (Eds.), Vol. 11693. Springer, Heidelberg, 269–295. https://doi.org/10.1007/978-3-030-26951-7_10
- [2] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. 2009. Fast Cryptographic Primitives and Circular-Secure Encryption Based on Hard Learning Problems. In *CRYPTO 2009 (LNCS)*, Shai Halevi (Ed.), Vol. 5677. Springer, Heidelberg, 595–618. https://doi.org/10.1007/978-3-642-03356-8_35
- [3] Sanjeev Arora and Rong Ge. 2011. New Algorithms for Learning in Presence of Errors. In *ICALP 2011, Part I (LNCS)*, Luca Aceto, Monika Henzinger, and Jiri Sgall (Eds.), Vol. 6755. Springer, Heidelberg, 403–415. https://doi.org/10.1007/978-3-642-22006-7_34
- [4] Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. 2020. Practical Product Proofs for Lattice Commitments. In *CRYPTO (2) (Lecture Notes in Computer Science)*, Vol. 12171. Springer, 470–499.
- [5] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. 1998. Quantum Lower Bounds by Polynomials. In *39th FOCS*. IEEE Computer Society Press, 352–361. <https://doi.org/10.1109/SFCS.1998.743485>
- [6] Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. 2000. Public-Key Encryption in a Multi-user Setting: Security Proofs and Improvements. In *EUROCRYPT 2000 (LNCS)*, Bart Preneel (Ed.), Vol. 1807. Springer, Heidelberg, 259–274. https://doi.org/10.1007/3-540-45539-6_18
- [7] Mihir Bellare and Gregory Neven. 2006. Multi-signatures in the plain public-Key model and a general forking lemma. In *ACM CCS 2006*, Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati (Eds.). ACM Press, 390–399. <https://doi.org/10.1145/1180405.1180453>
- [8] Daniel Bernstein. Accessed May 2021. eBACS: ECRYPT Benchmarking of Cryptographic Systems. (Accessed May 2021). <https://bench.cr.yp.to/results-kem.html>.
- [9] Daniel J. Bernstein. 2015. Multi-user Schnorr security, revisited. Cryptology ePrint Archive, Report 2015/996. (2015). <https://eprint.iacr.org/2015/996>.
- [10] Nina Bindel, Mike Hamburg, Kathrin Hövelmanns, Andreas Hülsing, and Edoardo Persichetti. 2019. Tighter Proofs of CCA Security in the Quantum Random Oracle Model. In *TCC 2019, Part II (LNCS)*, Dennis Hofheinz and Alon Rosen (Eds.), Vol. 11892. Springer, Heidelberg, 61–90. https://doi.org/10.1007/978-3-030-36033-7_3
- [11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. 2011. Random Oracles in a Quantum World. In *ASIACRYPT 2011 (LNCS)*, Dong Hoon Lee and Xiaoyun Wang (Eds.), Vol. 7073. Springer, Heidelberg, 41–69. https://doi.org/10.1007/978-3-642-25385-0_3
- [12] Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. 2018. CRYSTALS - Kyber: A CCA-Secure Module-Lattice-Based KEM. In *EuroS&P*. IEEE, 353–367.
- [13] Hao Chen, Kristin E. Lauter, and Katherine E. Stange. 2016. Security Considerations for Galois Non-dual RLWE Families. In *SAC 2016 (LNCS)*, Roberto Avanzi and Howard M. Heys (Eds.), Vol. 10532. Springer, Heidelberg, 443–462. https://doi.org/10.1007/978-3-319-69453-5_24
- [14] Jan-Pieter D’Anvers, Angshuman Karmakar, Sujoy Sinha Roy, and Frederik Vercauteren. 2018. Saber: Module-LWR Based Key Exchange, CPA-Secure Encryption and CCA-Secure KEM. In *AFRICACRYPT (Lecture Notes in Computer Science)*, Vol. 10831. Springer, 282–305.
- [15] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. 2021. Online-Extractability in the Quantum Random-Oracle Model. Cryptology ePrint Archive, Report 2021/280. (2021). <https://eprint.iacr.org/2021/280>.
- [16] Eiichiro Fujisaki and Tatsuaki Okamoto. 1999. Secure Integration of Asymmetric and Symmetric Encryption Schemes. In *CRYPTO’99 (LNCS)*, Michael J. Wiener (Ed.), Vol. 1666. Springer, Heidelberg, 537–554. https://doi.org/10.1007/3-540-48405-1_34
- [17] Eiichiro Fujisaki and Tatsuaki Okamoto. 2013. Secure Integration of Asymmetric and Symmetric Encryption Schemes. *Journal of Cryptology* 26, 1 (Jan. 2013), 80–101. <https://doi.org/10.1007/s00145-011-9114-1>
- [18] Mike Hamburg. 201. Post-quantum cryptography proposal: ThreeBears. <https://www.shiftright.org/papers/threebears/nist-submission.pdf>. (201).
- [19] Mike Hamburg. 2019. Security proof of ThreeBears. <https://www.shiftright.org/papers/threebears/proof-round2.pdf>. (2019).
- [20] Johan Håstad. 1988. Solving Simultaneous Modular Equations of Low Degree. *SIAM J. Comput.* 17, 2 (1988), 336–341. <https://doi.org/10.1137/0217019>
- [21] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. 2017. A Modular Analysis of the Fujisaki-Okamoto Transformation. In *TCC 2017, Part I (LNCS)*, Yael Kalai and Leonid Reyzin (Eds.), Vol. 10677. Springer, Heidelberg, 341–371. https://doi.org/10.1007/978-3-319-70500-2_12
- [22] Kathrin Hövelmanns. 2021. *Generic constructions of quantum-resistant cryptosystems*. doctoralthesis. Ruhr-Universität Bochum, Universitätsbibliothek. <https://doi.org/10.13154/294-7758>
- [23] Kathrin Hövelmanns, Eike Kiltz, Sven Schäge, and Dominique Unruh. 2020. Generic Authenticated Key Exchange in the Quantum Random Oracle Model. In *PKC 2020, Part II (LNCS)*, Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas (Eds.), Vol. 12111. Springer, Heidelberg, 389–422. https://doi.org/10.1007/978-3-030-45388-6_14
- [24] Haodong Jiang, Zhenfeng Zhang, Long Chen, Hong Wang, and Zhi Ma. 2018. IND-CCA-Secure Key Encapsulation Mechanism in the Quantum Random Oracle Model, Revisited. In *CRYPTO 2018, Part III (LNCS)*, Hovav Shacham and Alexandra Boldyreva (Eds.), Vol. 10993. Springer, Heidelberg, 96–125. https://doi.org/10.1007/978-3-319-96878-0_4
- [25] Haodong Jiang, Zhenfeng Zhang, and Zhi Ma. 2019. Key Encapsulation Mechanism with Explicit Rejection in the Quantum Random Oracle Model. In *PKC 2019*,

- Part II (LNCS), Dongdai Lin and Kazue Sako (Eds.), Vol. 11443. Springer, Heidelberg, 618–645. https://doi.org/10.1007/978-3-030-17259-6_21
- [26] Haodong Jiang, Zhenfeng Zhang, and Zhi Ma. 2019. Tighter Security Proofs for Generic Key Encapsulation Mechanism in the Quantum Random Oracle Model. In *Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019*, Jintai Ding and Rainer Steinwandt (Eds.), Springer, Heidelberg, 227–248. https://doi.org/10.1007/978-3-030-25510-7_13
- [27] Veronika Kuchta, Amin Sakzad, Damien Stehlé, Ron Steinfeld, and Shifeng Sun. 2020. Measure-Rewind-Measure: Tighter Quantum Random Oracle Model Proofs for One-Way to Hiding and CCA Security. In *EUROCRYPT 2020, Part III (LNCS)*, Anne Canteaut and Yuval Ishai (Eds.), Vol. 12107. Springer, Heidelberg, 703–728. https://doi.org/10.1007/978-3-030-45727-3_24
- [28] Adeline Langlois and Damien Stehlé. 2015. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.* 75, 3 (2015), 565–599.
- [29] Asheesh Laroia. 2011. Short key IDs are bad news (with OpenPGP and GNU Privacy Guard). <http://www.asheesh.org/note/debian/short-key-ids-are-bad-news.html>. (2011).
- [30] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. 2010. On Ideal Lattices and Learning with Errors over Rings. In *EUROCRYPT 2010 (LNCS)*, Henri Gilbert (Ed.), Vol. 6110. Springer, Heidelberg, 1–23. https://doi.org/10.1007/978-3-642-13190-5_1
- [31] NIST. 2016. Submission Requirements and Evaluation Criteria for the Post-Quantum Cryptography Standardization Process. <https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf>. (2016).
- [32] Oded Regev. 2005. On lattices, learning with errors, random linear codes, and cryptography. In *37th ACM STOC*, Harold N. Gabow and Ronald Fagin (Eds.), ACM Press, 84–93. <https://doi.org/10.1145/1060590.1060603>
- [33] Tsunekazu Saito, Keita Xagawa, and Takashi Yamakawa. 2018. Tightly-Secure Key-Encapsulation Mechanism in the Quantum Random Oracle Model. In *EUROCRYPT 2018, Part III (LNCS)*, Jesper Buus Nielsen and Vincent Rijmen (Eds.), Vol. 10822. Springer, Heidelberg, 520–551. https://doi.org/10.1007/978-3-319-78372-7_17

6 QUANTUM PRELIMINARIES

QUBITS. For simplicity, we will treat a *qubit* as a vector $|\varphi\rangle \in \mathbb{C}^2$, i.e., a linear combination $|\varphi\rangle = \alpha \cdot |0\rangle + \beta \cdot |1\rangle$ of the two *basis states* (vectors) $|0\rangle$ and $|1\rangle$ with the additional requirement to the probability amplitudes $\alpha, \beta \in \mathbb{C}$ that $|\alpha|^2 + |\beta|^2 = 1$. The basis $\{|0\rangle, |1\rangle\}$ is called *standard orthonormal computational basis*. The qubit $|\varphi\rangle$ is said to be *in superposition*. Classical bits can be interpreted as quantum bits via the mapping ($b \mapsto 1 \cdot |b\rangle + 0 \cdot |1 - b\rangle$).

QUANTUM REGISTERS. We will treat a quantum register as a collection of multiple qubits, i.e. a linear combination $|\varphi\rangle := \sum_{x \in \{0,1\}^n} \alpha_x \cdot |x\rangle$, where $\alpha_x \in \mathbb{C}$, with the additional restriction that $\sum_{x \in \{0,1\}^n} |\alpha_x|^2 = 1$. As in the one-dimensional case, we call the basis $\{|x\rangle\}_{x \in \{0,1\}^n}$ the *standard orthonormal computational basis*. We say that $|\varphi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x \cdot |x\rangle$ contains the classical query x if $\alpha_x \neq 0$.

MEASUREMENTS. Qubits can be measured with respect to a basis. In this paper, we will only consider measurements in the standard orthonormal computational basis, and denote this measurement by $\text{MEASURE}(\cdot)$, where the outcome of $\text{MEASURE}(|\varphi\rangle)$ for a single qubit $|\varphi\rangle = \alpha \cdot |0\rangle + \beta \cdot |1\rangle$ will be 0 with probability $|\alpha|^2$ and 1 with probability $|\beta|^2$, and the outcome of measuring a qubit register $|\varphi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x \cdot |x\rangle$ will be x with probability $|\alpha_x|^2$. Note that the amplitudes *collapse* during a measurement, this means that by measuring $\alpha \cdot |0\rangle + \beta \cdot |1\rangle$, α and β are switched to one of the combinations in $\{\pm(1, 0), \pm(0, 1)\}$. Likewise, in the n -dimensional case, all amplitudes are switched to 0 except for the one that belongs to the measurement outcome and which will be switched to 1.

QUANTUM ORACLES AND QUANTUM ADVERSARIES. Following [5, 11], we view a quantum oracle $|O\rangle$ as a mapping

$$|x\rangle|y\rangle \mapsto |x\rangle|y \oplus O(x)\rangle,$$

where $O : \{0, 1\}^n \rightarrow \{0, 1\}^m$, and model quantum adversaries \mathcal{A} with access to O by a sequence $U_1, |O\rangle, U_2, \dots, |O\rangle, U_N$ of unitary transformations. We write $\mathcal{A}^{(O)}$ to indicate that the oracles are quantum-accessible (contrary to oracles which can only process classical bits).

QUANTUM RANDOM ORACLE MODEL. We consider security games in the quantum random oracle model (QROM) as their counterparts in the classical random oracle model, with the difference that we consider quantum adversaries that are given *quantum* access to the (offline) random oracles involved, and *classical* access to all other (online) oracles.

7 PROOF OF THEOREM 3.1 (PART 2)

PROOF. Let \mathcal{A} be an adversary and consider the games given in Fig. 10.

GAME G_0 . This is the original (n, q_C) -IND-CCA game. Thus,

$$\left| \Pr \left[G_0^{\mathcal{A}} \Rightarrow 1 \right] - \frac{1}{2} \right| = \text{Adv}_{\text{KEM}}^{(n, q_C)\text{-IND-CCA}}(\mathcal{A}).$$

GAME G_1 . In game G_1 the challenge oracle samples m from $\mathcal{M} \setminus \mathcal{L}_M$ instead of \mathcal{M} , since on a message collision an adversary can easily distinguish the real KEM key from the random one. By the birthday bound,

$$\left| \Pr \left[G_0^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[G_1^{\mathcal{A}} \Rightarrow 1 \right] \right| \leq \frac{q_C^2}{|\mathcal{M}|}.$$

GAME G_2 . In game G_2 we modify the Decaps oracle in lines 21 and 22 such that for an invalid ciphertext the key is defined as $K'_j(c)$, where K'_j is an independent internal random oracle. This remains unnoticed to adversary \mathcal{A} unless it queries

$$K(s_j, \cdot),$$

for $j \in [1, n]$ or there exists a collision $s_i = s_j$, that is there exists $i, j \in [n]$ s.t. $s_i = s_j$, in which case the event BAD is set to true.

Since the seeds $s_j \in \{0, 1\}^\lambda$ are uniformly random and information-theoretically hidden from the adversary, we have by the birthday and union bound

$$\left| \Pr \left[G_2^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[G_1^{\mathcal{A}} \Rightarrow 1 \right] \right| \leq \Pr[\text{BAD}] \leq \frac{n^2 + n \cdot q_F}{2^\lambda}.$$

GAME G_3 In game G_3 we start introducing changes to the decapsulation oracle $\text{Decaps}(j, \cdot)$ with the goal of simulating decapsulation queries without the secret-keys by additional changes introduced in game G_4 . In more detail, for $m' := \text{Dec}(sk_j, c)$ we return the “implicit reject key” $K'_j(c)$ if m' was not queried on F , and otherwise return the key K from $(K, r) \leftarrow F(m')$. In other words, we consider only ciphertexts as valid, where the adversary first queries G (indirectly through F). The probability that a ciphertext is valid without first querying G is bounded by the γ -spreadness. Note that this step is necessary due to not domain separating the random oracle H in order to guarantee the consistency of the decapsulation query lists \mathcal{L}_D and the random oracle list \mathcal{L}_H . To see why this is necessary, consider two ciphertexts c_1 and c_2 which were not derived by making queries first to G (indirectly through F). A priori,

GAMES $G_0 - G_5$	F(A)
01 for $j \in [n]$	34 if $A \in \mathcal{M}$
02 $(pk_j, sk_j) \xleftarrow{\$} \text{Gen}$	35 $m := A$
03 $s_j \xleftarrow{\$} \{0, 1\}^\lambda$	36 $(K, r) := (H(m), G(m))$
04 if $\exists i, j \in [n]$ s.t. $s_i = s_j$; BAD := true ; abort	37 if $A \in \{0, 1\}^\lambda \times \mathcal{C}$
05 $\vec{pk} \leftarrow (pk_1, \dots, pk_n)$	38 $(s, c) := A$
06 $b \xleftarrow{\$} \{0, 1\}$	39 $K := K(s, c)$
07 $b' \xleftarrow{\$} \mathcal{A}^{\text{Decaps, Chall, F}}(\vec{pk})$	40 if K undefined: $K \xleftarrow{\$} \mathcal{K}$
08 return $\llbracket b' = b \rrbracket$	41 if r undefined: $r \xleftarrow{\$} \mathcal{R}$
	42 return (K, r)
Chall(j)	G(m) / Internal random oracle
09 $m \xleftarrow{\$} \mathcal{M}$	43 $r \xleftarrow{\$} \mathcal{R}$
10 $m \xleftarrow{\$} \mathcal{M} \setminus \mathcal{L}_M$	44 if $m \in \mathcal{L}_M$ / G_5
11 $(K_0, r) \leftarrow F(m)$	45 QUERY := true / G_5
12 $\mathcal{L}_M := \mathcal{L}_M \cup \{m\}$	46 abort / G_5
13 $r \xleftarrow{\$} \mathcal{R}$; $K_0 \xleftarrow{\$} \{0, 1\}^k$	47 return r
14 $K_1 \xleftarrow{\$} \{0, 1\}^k$	
15 $c \leftarrow \text{Enc}(pk_j, m; r)$	K(s, c) / Internal random oracle
16 $\mathcal{L}_{C_j} = \mathcal{L}_{C_j} \cup \{c\}$	48 $K \xleftarrow{\$} \mathcal{K}$
17 return (c, K_b)	49 if $s \in \{s_1, \dots, s_n\}$ / G_2-G_5
Decaps(j, c \notin \mathcal{L}_{C_j})	50 BAD := true / G_2-G_5
18 $m' := \text{Dec}(sk_j, c)$	51 abort / G_2-G_5
19 $(K, r) \leftarrow F(m')$	52 return K
20 if $m' = \perp$ or $\text{Enc}(pk_j, m'; r) \neq c$	H(m) / Internal random oracle
21 $(K, r) \leftarrow F(s_j, c)$	53 $K \xleftarrow{\$} \mathcal{K}$
22 $K := K'_j(c)$	54 if $m \in \mathcal{L}_M$ / G_5
23 return K	55 QUERY := true / G_5
Decaps(j, c \notin \mathcal{L}_{C_j})	56 abort / G_5
24 $m' := \text{Dec}(sk_j, c)$	57 for $j \in [n]$: / G_4-G_5
25 if $\exists (m, r, K) \in \mathcal{L}_F$ with $\text{Enc}(pk_j, m; r) = c$ and $m = m'$	58 $c'_j := \text{Enc}(pk_j, m; G(m))$ / G_4-G_5
26 return K	59 $\mathcal{L}_{D_j} := \mathcal{L}_{D_j} \cup \{(c'_j, K)\}$ / G_4-G_5
27 $K := K'_j(c)$	60 return K
28 return K	
Decaps(j, c \notin \mathcal{L}_{C_j})	
29 if $\exists K$ s. th. $(c, K) \in \mathcal{L}_{D_j}$ / G_4-G_5	
30 return K	
31 $K \xleftarrow{\$} \mathcal{K}$	
32 $\mathcal{L}_{D_j} := \mathcal{L}_{D_j} \cup \{(c, K)\}$	
33 return K	

Figure 10: Games $G_0 - G_4$ for the proof of Theorem 3.1. The internal random oracles H'_j, H_j, G_j, K_j are not accessible by the adversary. We assume wlog that F is only queried once on each value A .

we can not exclude the possibility that c_1 under secret-key sk_1 decrypts to the same m as c_2 under the secret-key sk_2 . With this game hop what we are essentially showing is that this does not matter, since by γ -spreadness, the re-encryption check will not pass, and thus the KEM keys will be independent. Without this argument, the simulation would return different KEM keys although we have not excluded that re-encryption might be successful, where the simulation would then need to return the same KEM key, i.e. $H(m)$. By the weak γ -spreadness we have by a union bound over the n

users and q_D decapsulation queries

$$\left| \Pr \left[G_2^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[G_3^{\mathcal{A}} \Rightarrow 1 \right] \right| \leq q_D n 2^{-\gamma}.$$

GAME G_4 . In game G_4 we simulate the decapsulation oracle $\text{Decaps}(j, \cdot)$ without knowledge of the secret key by patching the random oracle H in lines 57-59 and 29-33. Note that if PKE was perfectly correct, then the random oracle patching is also perfectly correct and therefore two games would look identical in \mathcal{A} 's view.

The only bad case happens if G is queried on some m which induces a correctness error, that is $\text{Dec}(sk_j, \text{Enc}(pk_j; m; G(m))) \neq$

m . More concretely, define the sets $\text{BAD}_j := \{m \in \mathcal{M} \mid m \neq m', \text{ where } c \leftarrow \text{Enc}(pk_j; m; G(m)); m' \leftarrow \text{Dec}(sk_j, c)\}$.

Define the event CORR_j to be the event that \mathcal{A} makes an (implicit) query to $G(m)$ for some $m \in \text{BAD}_j$. Since there are at most $(q_F + q_D)$ explicit and implicit queries to G , we have

$$\Pr[\text{CORR}_j] \leq (q_F + q_D)\delta.$$

Defining $\text{CORR} := \bigcup_{j=1}^n \text{CORR}_j$, we obtain by the union bound⁴

$$\Pr[\text{CORR}] \leq n(q_F + q_D)\delta.$$

We now claim

$$\left| \Pr[G_4^{\mathcal{A}} \Rightarrow 1] - \Pr[G_3^{\mathcal{A}} \Rightarrow 1] \right| \leq \Pr[\text{CORR}].$$

Let us analyze why G_4 and G_3 are identical conditioned on $\neg\text{CORR}$. Consider a query $\text{Decaps}(j, c)$ and define $m'_j := \text{Dec}(sk_j, c)$ and $c'_j := \text{Enc}(pk_j, m'_j; G(m'_j))$.

- **Case 1:** $m'_j = \perp$. H cannot be called on $m'_j = \perp$ and hence the KEM key of $K = \text{Decaps}(sk_j, c) = K'_j(c)$ in G_3 is identically distributed as $(c, K) \in \mathcal{L}_{D_j}$ in G_4 .
- **Case 2:** $m'_j \neq \perp \wedge c \neq c'_j$. Both games return a uniform random key K . The only way for \mathcal{A} to detect a difference between the two games is if it makes a query $H(m)$ such that $\text{Enc}(pk_j, m; G(m)) = c$. But conditioned that $m \notin \text{BAD}_j$, c does not lie in the image space of encryption $\text{Im}(\text{Enc}_{pk_j}(-; G(-)))$, since the re-encryption check failed and therefore such an m cannot exist.
- **Case 3:** $m'_j \neq \perp \wedge c = c'_j$. In game G_3 , $\text{Decaps}(j, c)$ returns $K = H(m'_j)$, whereas G_4 picks a uniform K in $H(m)$ to match K for all m that deterministically encrypt to the same c , i.e., all m satisfying $\text{Enc}(pk_j, m; G(m)) = c$ and saves those values in \mathcal{L}_{D_j} . Since $c = c'_j$ the re-encryption check had to pass and H was queried before Decaps . The only way for \mathcal{A} to detect a difference between the two games is to query H on some value $m \neq m'_j$ that also deterministically encrypts to the same c . But this also implies that $G(m)$ was queried for some $m \in \text{BAD}_j$.

In order to see why we had to use γ -spreadness, consider $\text{Decaps}(i, c)$ and define $m'_i := \text{Dec}(sk_i, c_i)$ and $c'_i := \text{Enc}(pk_i, m'_i; G(m'_i))$. Then, there are two additional cases to consider, since the random oracle H is shared between the users.

- **Case 1:** $m'_j = m'_i$. Then either both c and c_i are valid ciphertexts or both of them are invalid, i.e. do not pass the re-encryption check. If the adversary first queries F (and implicitly H) first and then Decaps the oracles of G_3 and G_4 are consistent, due to the patching in lines 57-59. If the adversary queries first Decaps on c and c_i and then H , they will be invalid ciphertexts and therefore both G_3 and G_4 correctly give out independent “implicit rejection” KEM keys (line 27 in G_3 and line 31 in G_4). This case is the reason we introduced the game hop from G_2 to G_3 where we applied γ -spreadness.
- **Case 2:** $m'_j \neq m'_i$. Here, there are no new potential inconsistency problems due to the multi-user setting, which might occur since $H(m'_j)$ and $H(m'_i)$ are independent.

GAME G_5 . In game G_5 we abort on queries of the form $H(m)$ or $G(m)$ for some challenge message $m \in \mathcal{L}_M$, in which case the event QUERY holds true. We have by the difference lemma,

$$\left| \Pr[G_5^{\mathcal{A}} \Rightarrow 1] - \Pr[G_4^{\mathcal{A}} \Rightarrow 1] \right| \leq \Pr[\text{QUERY}].$$

Note that in G_5 bit b is independent of the view of the adversary. We thus have

$$\Pr[G_5^{\mathcal{A}} \Rightarrow 1] = \frac{1}{2}.$$

Analogously to (3), we can again prove that there exists an adversary C with

$$\Pr[\text{QUERY}] \leq 2 \cdot \left(\text{Adv}_{\text{PKE}}^{(n, q_C)\text{-IND-CPA}}(C) + \frac{q_F q_C}{|\mathcal{M}|} \right).$$

Summing up the inequalities yields the claimed bound, concluding the proof of the theorem. \square

⁴Technically, we would need to consider the δ -correctness for fixed keys pk, sk and then take the expectations to be formally correct. We omit this for brevity.