

Tighter proofs of CCA security in the quantum random oracle model

Nina Bindel¹, Mike Hamburg², Andreas Hülsing³, and Edoardo Persichetti⁴

¹ Technische Universität Darmstadt, nbindel@cdc.informatik.tu-darmstadt.de

² Rambus, mhamburg@rambus.com

³ Eindhoven University of Technology, andreas@huelising.net

⁴ Florida Atlantic University, epersichetti@fau.edu

Abstract. We revisit the construction of IND-CCA secure key encapsulation mechanisms (KEM) from public-key encryption schemes (PKE). We give new, tighter security reductions for several constructions. Our main result is a tight reduction for the security of the U^\perp -transform of Hofheinz, Hövelmanns, and Kiltz (TCC'17) which turns OW-CPA secure deterministic PKEs into IND-CCA secure KEMs. This result is enabled by a new one-way to hiding (O2H) lemma which gives a tighter bound than previous O2H lemmas in certain settings and might be of independent interest. We extend this result also to the case of PKEs with non-zero decryption failure probability, partially non-injective PKEs, and non-deterministic PKEs.

In addition, we analyze the impact of different variations of the U^\perp -transform discussed in the literature on the security of the final scheme. We consider the difference between explicit (U^\perp) and implicit (U^\perp) rejection, proving that security of the former implies security of the latter. We show that the opposite direction holds if the scheme with explicit rejection also uses key confirmation. Finally, we prove that (at least from a theoretic point of view) security is independent of whether the session keys are derived from message and ciphertext (U^\perp) or just from the message (U_m^\perp).

1 Introduction

If a general-purpose quantum computer can be built, it will break most widely-deployed public-key cryptography. The cryptographic community is busily designing new cryptographic systems to prepare for this risk. These systems typically consist of an algebraic structure with cryptographic hardness properties, plus a symmetric cryptography layer which transforms the algebraic structure into a higher level primitive like a public-key encryption (PKE) scheme, a key encapsulation mechanism (KEM), or a signature scheme. The algebraic structures underlying these so-called “post-quantum” systems have new properties, and the quantum threat model requires changes in the way security is analyzed. Therefore the transformations turning the algebraic structures into cryptosystems have to be freshly examined.

In this work we focus on the construction of secure KEMs. In this setting the algebraic structures usually provide a PKE from which a KEM is derived via a generic transform. A new property of the algebraic structures used in many post-quantum PKEs and KEMs gives them malleable ciphertexts, so they are at risk from chosen-ciphertext attacks (CCA) [HNP⁺03]. The standard defenses against CCA are variants of the Fujisaki-Okamoto (FO) transform [FO99]. Known security proofs for the FO transform use the random oracle model (ROM) [BR93]. This is for two reasons. First, the FO transform has a circular structure—it chooses coins for encryption according to the message being encrypted. This leads to obstacles which we do not know how to overcome when proving security in the standard model. In the ROM, we circumvent this by re-programming. Second, in the ROM a reduction learns all the adversary’s queries to the random oracle. This allows us to formalize the intuition that an adversary must have known a challenge plaintext and to extract said plaintext.

Since we are concerned with security against quantum attackers, we need to extend these proofs to the quantum-accessible random oracle model (QROM) [BDF⁺11]. This comes with two challenges for our setting. On the one hand, in the QROM the adversary can query all inputs in superposition. Hence, it is no longer trivial to break the circular dependency by re-programming, which results in security bounds that do not tightly match known attacks. On the other hand, a reduction cannot learn the adversarial queries by simple observation anymore. The reason is that observation of a quantum state requires a measurement which disturbs the state. Hence, more advanced techniques are required.

1.1 Our Contribution

QROM analysis of KEMs has advanced rapidly over the past several years. The initial solutions were loose by a factor of up to q^6 [TU16,HHK17], where q is the number of times the adversary queries the random oracle. This has improved to q^2 [SXY18,JZC⁺17] and finally to q [JZM19a,JZM19c]. Some work provides tight proofs under stronger assumptions [SXY18,XY18]. Our work provides a tight proof of KEM security for KEMs constructed from deterministic PKEs (Theorem 2), except for a square-root loss which might be impossible to avoid [JZM19b]. For KEMs constructed from randomized PKEs our bound is still loose by a factor of up to q (Theorem 1).

At the heart of our bound is a new one-way to hiding (O2H) lemma which gives a tighter bound than previous O2H lemmas (Lemma 5). This comes at the cost of limited applicability. O2H lemmas allow to bound the difference in the success probability of an adversary when replacing its oracle function by a similar function. Previous lemmas lost a factor of roughly the number of the adversary’s queries to this oracle or its square-root. Our lemma does not incur any such loss. On the downside, our lemma only applies if the reduction has access to both oracle functions and if the functions only differ in one position. See Table 1 for a comparison.

Some post-quantum schemes feature an inherent probability of decryption failure, say $\delta > 0$. Such failures can be used in attacks, but they also complicate

security proofs. As a result, previous bounds typically contain a term $q\sqrt{\delta}$ which is not known to be tight. However, most of the obstacles that arise in our CCA security proof can be avoided by assuming that encryption with a particular public key is injective. This is generally the case, even for imperfectly-correct systems; see Appendix D for a rough analysis of LWE schemes. In that case, the adversary’s advantage is limited to the probability that it actually finds and submits a valid message that fails to decrypt. This means that our bounds apply to deterministic but failure-prone systems like one of the BIKE [ABB⁺19] variants (BIKE-2).

Until today several variants of the FO-transform were proposed. We consider the four basic transforms $U^\perp, U_m^\perp, U^\times, U_m^\times$ and, in addition, we study U_m^\perp in the presence of key confirmation. The two most notable differences reside in the use of implicit rejection (U^\times, U_m^\times) versus explicit rejection (U^\perp, U_m^\perp), and whether the derivation of the session key should depend on the ciphertext (U_m^\perp, U_m^\times) or not (U^\perp, U^\times). Another important decision is the use of key confirmation which we also partially analyze. We come to the following results. Security with implicit rejection implies security with explicit rejection (Theorem 3). The opposite holds if the scheme with explicit rejection also employs key confirmation (Theorem 4). Moreover, security is independent of the decision if the session key derivation depends on the ciphertext (Theorem 5).

Notation. We will use the following notation throughout the paper.

- For two sets X, Y , we write Y^X to denote the set of functions from X to Y .
- Let $H : X \rightarrow Y$ be a (classical or quantum-accessible) random oracle. Then we denote the programming of H at $x \in X$ to some $y \in Y$ as $H[x \rightarrow y]$.
- Let \mathcal{A} be an algorithm. If \mathcal{A} has access to a classical (resp., quantum-accessible) oracle H , we write \mathcal{A}^H and call \mathcal{A} an oracle (resp., quantum oracle) algorithm.

2 One-way to Hiding

ROM reductions typically simulate the random oracle in order to learn the adversary’s queries. In the classical ROM, the adversary cannot learn any information about $H(x)$ without the simulator learning both x and $H(x)$. In the QROM things are not so simple, because measuring or otherwise recording the queries might collapse the adversary’s quantum state and change its behavior. However, under certain conditions the simulator can learn the queries using “One-way to Hiding” (O2H) techniques going back to [Unr15]. We will use the O2H techniques from [AHU18], and introduce a novel variant that allows for tighter results.

Consider two quantum-accessible oracles $G, H : X \rightarrow Y$. The oracles do not need to be random. Suppose that G and H differ only on some small set $S \subset X$, meaning that $\forall x \notin S, G(x) = H(x)$. Let \mathcal{A} be an oracle algorithm that takes an input z and makes at most q queries to G or H . Possibly \mathcal{A} makes them in parallel. Therefore, suppose that the query depth, i.e., the maximum number of sequential invocations of the oracle [AHU18], is at most $d \leq q$. If $\mathcal{A}^G(z)$ behaves

differently from $\mathcal{A}^H(z)$, then the O2H techniques give a way for the simulator to find some $x \in S$ with probability dependent on d and q .

We will use the following three O2H lemmas.

- Lemma 1 (original O2H) is the most general: the simulator needs to provide only G or H but it has the least probability of success.
- Lemma 3 (semiclassical O2H) has a greater probability of success, but requires more from the simulator: for each query x , the simulator must be able to recognize whether $x \in S$, and if not it must return $G(x) = H(x)$.
- Lemma 5 (our new “double-sided” O2H) gives the best probability of success, but it requires the simulator to evaluate both G and H in superposition. It also can only extract $x \in S$ if S has a single element. If S has many elements, but the simulator knows a function f such that $\{f(x) : x \in S\}$ has a single element, then it can instead extract that element $f(x)$.

We summarize the three variants of O2H as shown in Table 1. In all cases, there are two oracles H and G that differ in some set S , and the simulator outputs $x \in S$ with some probability ϵ . The lemma then shows an upper bound on the difference between \mathcal{A}^H and \mathcal{A}^G as a function of ϵ .

Variant	Lemma	Ref	Oracles differ	Sim. must know	Bound
Original	Lem. 1	[AHU18]	Arbitrary	H or G	$2d\sqrt{\epsilon}$
Semi-classical	Lem. 3	[AHU18]	Arbitrary	S and $(H \setminus S$ or $G \setminus S)$	$2\sqrt{d\epsilon}$
Double-sided	Lem. 5	this work	One place	H and G	$2\sqrt{\epsilon}$

Table 1. Comparison of O2H variants

Arbitrary joint distribution The O2H lemmas allow (G, H, S, z) to be random with arbitrary joint distribution. This is stronger than (G, H, S, z) being arbitrary fixed objects, because the probabilities in the lemma include the choice of (G, H, S, z) in addition to \mathcal{A} 's coins and measurements. Also, the lemmas are still true if the adversary consults other oracles which are also drawn from a joint distribution with (G, H, S, z) .

2.1 Original O2H

We begin with the original O2H which first appeared in [Unr15]. We use the phrasing from [AHU18] as it is more general and more consistent with our other lemmata.

Lemma 1 (One-way to hiding; [AHU18] Theorem 3). *Let $G, H : X \rightarrow Y$ be random functions, let z be a random value, and let $S \subset X$ be a random set such that $\forall x \notin S, G(x) = H(x)$. (G, H, S, z) may have arbitrary joint distribution. Furthermore, let \mathcal{A}^H be a quantum oracle algorithm which queries H with depth at most d . Let Ev be an arbitrary classical event. Define an oracle algorithm $\mathcal{B}^H(z)$ as follows: Pick $i \xleftarrow{\$} \{1, \dots, d\}$. Run $\mathcal{A}^H(z)$ until just before its i th round*

of queries to H . Measure all query input registers in the computational basis, and output the set T of measurement outcomes. Let

$$P_{\text{left}} := \Pr[\text{Ev} : \mathcal{A}^H(z)], \quad P_{\text{right}} := \Pr[\text{Ev} : \mathcal{A}^G(z)],$$

$$P_{\text{guess}} := \Pr[S \cap T \neq \emptyset : T \leftarrow \mathcal{B}^H(z)].$$

Then

$$|P_{\text{left}} - P_{\text{right}}| \leq 2d\sqrt{P_{\text{guess}}} \quad \text{and} \quad \left| \sqrt{P_{\text{left}}} - \sqrt{P_{\text{right}}} \right| \leq 2d\sqrt{P_{\text{guess}}}.$$

The same result holds with $\mathcal{B}^G(z)$ instead of $\mathcal{B}^H(z)$ in the definition of P_{guess} .

From this lemma we conclude the following result for pseudo-random functions (PRFs, see Definition 10). It intuitively states that a random oracle makes a good PRF, even if the distinguisher is given full access to the random oracle in addition to the PRF oracle.

Corollary 1 (PRF based on random oracle). *Let $H : (K \times X) \rightarrow Y$ be a quantum-accessible random oracle. This function may be used as a quantum-accessible PRF $F_k(x) := H(k, x)$ with a key $k \xleftarrow{\$} K$. Suppose a PRF-adversary \mathcal{A} makes q queries to H at depth d , and any number of queries to F_k at any depth. Then*

$$\text{Adv}_{F_k}^{\text{PRF}}(\mathcal{A}) \leq 2\sqrt{dq/|K|}.$$

Proof. The adversary's goal is to distinguish (F_k, H) from (F, H) , where F is an unrelated uniformly random function. This is the same as distinguishing $(F, H[(k, x) \rightarrow F(x)])$ from (F, H) , and the set of differences between these two H -oracles is $S := \{k\} \times X$. By Lemma 1, the distinguishing advantage is at most $2d\sqrt{P_{\text{guess}}}$, where $P_{\text{guess}} = \Pr[\exists(k', x) \in Q : k' = k]$, for a random round Q of parallel queries made by $\mathcal{A}^{F, H}$.

Since $\mathcal{A}^{F, H}$ has no information about k , and in expectation Q contains q/d parallel queries, we have $P_{\text{guess}} \leq q/(d \cdot |K|)$, so

$$\text{Adv}_{F_k}^{\text{PRF}}(\mathcal{A}) \leq 2d\sqrt{q/(d \cdot |K|)} = 2\sqrt{dq/|K|}$$

as claimed. □

2.2 Semi-classical O2H

We now move on to semi-classical O2H. Here \mathcal{B} is defined in terms of *punctured oracles* [AHU18], which measure whether the input is in a set S as defined next.

Definition 1 (Punctured oracle). *Let $H : X \rightarrow Y$ be any function, and $S \subset X$ be a set. The oracle $H \setminus S$ (“ H punctured by S ”) takes as input a value x . It first computes whether $x \in S$ into an auxiliary qubit p , and measures p . Then it runs $H(x)$ and returns the result. Let Find be the event that any of the measurements of p returns 1.*

The event is called Find because if the simulator chooses to, it can immediately terminate the simulation and measure the value $x \in S$ which caused the event. The oracle is called “punctured” because if Find does not occur, $H \setminus S$ returns a result independent of H ’s outputs on S , as shown by the following lemma.

Lemma 2 (Puncturing is effective; [AHU18] Lemma 1). *Let $G, H : X \rightarrow Y$ be random functions, let z be a random value, and let $S \subset X$ be a random set such that $\forall x \notin S, G(x) = H(x)$. (G, H, S, z) may have arbitrary joint distribution. Let \mathcal{A}^H be a quantum oracle algorithm. Let Ev be an arbitrary classical event. Then*

$$\Pr[\text{Ev} \wedge \neg \text{Find} : x \leftarrow \mathcal{A}^{H \setminus S}(z)] = \Pr[\text{Ev} \wedge \neg \text{Find} : x \leftarrow \mathcal{A}^{G \setminus S}(z)].$$

Also, puncturing only disturbs the adversary’s state when it is likely to Find.

Lemma 3 (Semi-classical O2H; [AHU18] Theorem 1). *Let $G, H : X \rightarrow Y$ be random functions, let z be a random value, and let $S \subset X$ be a random set such that $\forall x \notin S, G(x) = H(x)$. (G, H, S, z) may have arbitrary joint distribution.*

Let \mathcal{A}^H be a quantum oracle algorithm which queries H with depth at most d . Let Ev be an arbitrary classical event and let

$$\begin{aligned} P_{\text{left}} &:= \Pr[\text{Ev} : \mathcal{A}^H(z)], & P_{\text{right}} &:= \Pr[\text{Ev} : \mathcal{A}^G(z)], \\ P_{\text{find}} &:= \Pr[\text{Find} : \mathcal{A}^{G \setminus S}(z)] & \stackrel{\text{Lem. 2}}{=} & \Pr[\text{Find} : \mathcal{A}^{H \setminus S}(z)]. \end{aligned}$$

Then

$$|P_{\text{left}} - P_{\text{right}}| \leq 2\sqrt{dP_{\text{find}}} \quad \text{and} \quad \left| \sqrt{P_{\text{left}}} - \sqrt{P_{\text{right}}} \right| \leq 2\sqrt{dP_{\text{find}}}.$$

The theorem also holds with bound $\sqrt{(d+1)P_{\text{find}}}$ for the following alternative definitions of P_{right} :

$$\begin{aligned} P_{\text{right}} &:= \Pr[\text{Ev} : \mathcal{A}^{H \setminus S}(z)] \\ P_{\text{right}} &:= \Pr[\text{Ev} \wedge \neg \text{Find} : \mathcal{A}^{H \setminus S}(z)] & \stackrel{\text{Lem. 2}}{=} & \Pr[\text{Ev} \wedge \neg \text{Find} : \mathcal{A}^{G \setminus S}(z)] \\ P_{\text{right}} &:= \Pr[\text{Ev} \vee \text{Find} : \mathcal{A}^{H \setminus S}(z)] & \stackrel{\text{Lem. 2}}{=} & \Pr[\text{Ev} \vee \text{Find} : \mathcal{A}^{G \setminus S}(z)] \end{aligned}$$

We might expect that if the adversary has no information about S , then P_{find} would be at most $q|S|/|X|$. But this is not quite true: the disturbance caused by puncturing gives the adversary information about S . This increases \mathcal{A} ’s chances, but only by a factor of 4, as explained next.

Lemma 4 (Search in semi-classical oracle; [AHU18] Theorem 2). *Let $H : X \rightarrow Y$ be a random function, let z be a random value, and let $S \subset X$ be a random set. (H, S, z) may have arbitrary joint distribution. Let \mathcal{A}^H be a*

quantum oracle algorithm which queries H at most q times with depth at most d .

Let $\mathcal{B}^H(z)$ and P_{guess} be defined as in Lemma 1. Then

$$\Pr[\text{Find} : \mathcal{A}^{H \setminus S}(z)] \leq 4dP_{\text{guess}}.$$

In particular, if for each $x \in X$, $\Pr[x \in S] \leq \epsilon$ (conditioned on z , on other oracles \mathcal{A} has access to, and on other outputs of H) then

$$\Pr[\text{Find} : \mathcal{A}^{H \setminus S}(z)] \leq 4q\epsilon.$$

2.3 Double-sided O2H

We augment these lemmas with a new O2H lemma which achieves a tighter bound focusing on a special case. This focus comes at the price of limited applicability. Our lemma applies when the simulator can simulate both G and H .

Lemma 5 (Double-sided O2H). *Let $G, H : X \rightarrow Y$ be random functions, let z be a random value, and let $S \subset X$ be a random set such that $\forall x \notin S, G(x) = H(x)$. (G, H, S, z) may have arbitrary joint distribution. Let $f : X \rightarrow W \subseteq \{0, 1\}^n$ be a function such that*

$$\{f(x) : x \in S\} = \{w^*\}$$

is a single element. Let \mathcal{A}^H be a quantum oracle algorithm. Let Ev be an arbitrary classical event. We will define another quantum oracle algorithm $\mathcal{B}^{G,H}(z)$ which returns either w^ or a failure symbol \perp . \mathcal{B} runs in about the same amount of time as \mathcal{A} , but when \mathcal{A} queries H , \mathcal{B} queries both G and H , and also runs f twice. Let*

$$P_{\text{left}} := \Pr[\text{Ev} : \mathcal{A}^H(z)], \quad P_{\text{right}} := \Pr[\text{Ev} : \mathcal{A}^G(z)], \quad P_{\text{extract}} := \Pr[\mathcal{B}^{G,H}(z) \rightarrow w^*].$$

Then

$$|P_{\text{left}} - P_{\text{right}}| \leq 2\sqrt{P_{\text{extract}}} \quad \text{and} \quad \left| \sqrt{P_{\text{left}}} - \sqrt{P_{\text{right}}} \right| \leq 2\sqrt{P_{\text{extract}}}.$$

Proof. See Appendix B.

Note that if $S = \{x^*\}$ is already a single element, then we may take f as the identity. In this case \mathcal{B} will return either \perp or x^* .

3 KEM and PKE Security Proofs

We are now ready to get to the core of our work. All the relevant security notions are given in Appendix A. The implications are summarized in Figure 1.

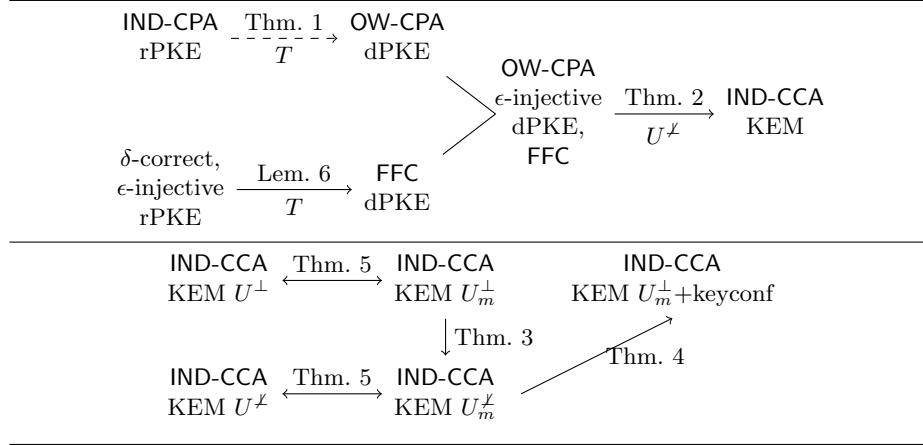


Fig. 1. Relations of our security notions using transforms T and U^\perp (above) and relations between the security of different types of U -constructions (below). The solid lines show implications which are tight with respect to powers of q and/or d , and the dashed line shows a non-tight implication.

3.1 Derandomization: $\text{IND-CPA } \mathbf{P} \xrightarrow{\text{QROM}} \text{OW-CPA } T(\mathbf{P}, G)$

The T transform [HHK17] converts a rPKE $\mathbf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ to a dPKE $T(\mathbf{P}, G) = (\text{Keygen}, \text{Encr}_1, \text{Decr})$ by using a hash function $G : \mathcal{M} \rightarrow \mathcal{R}$, modeled as random oracle, to choose encryption coins, where

$$\text{Encr}_1(\text{pk}, m) := \text{Encr}(\text{pk}, m; G(m)).$$

The following theorem shows that if a PKE \mathbf{P} is IND-CPA secure⁵, then $T(\mathbf{P}, G)$ is one-way secure in the quantum-accessible random oracle model.

Theorem 1. *Let \mathbf{P} be an rPKE with messages in \mathcal{M} and random coins in \mathcal{R} . Let $G : \mathcal{M} \rightarrow \mathcal{R}$ be a quantum-accessible random oracle. Let \mathcal{A} be an OW-CPA adversary against $\mathbf{P}' := T(\mathbf{P}, G)$. Suppose that \mathcal{A} queries G at most q times with depth at most d .*

Then we can construct an IND-CPA adversary \mathcal{B} against \mathbf{P} , running in about the same time and resources as \mathcal{A} , such that

$$\text{Adv}_{\mathbf{P}'}^{\text{OW-CPA}}(\mathcal{A}) \leq (d+2) \cdot \left(\text{Adv}_{\mathbf{P}}^{\text{IND-CPA}}(\mathcal{B}) + \frac{8(q+1)}{|\mathcal{M}|} \right).$$

Proof. See Appendix C.

Second preimages. In the traditional definition of one-way functions, the adversary wins by finding any m' where $\text{Encr}(\text{pk}, m') = c^*$, whereas in our definition

⁵ The theorem actually only requires a weaker notion, IND-KPA-security, in which the challenge messages are chosen at random instead of adversarially.

of OW-CPA the adversary must find m^* itself. This only matters if there is a second preimage, and thus a decryption failure. If P is δ -correct and ϵ -injective, it is easily shown that a definition allowing second preimages adds at most $\min(\delta, \epsilon)$ to the adversary’s OW-CPA-advantage.

Hashing the public key. Many KEMs use a variant of T which sets the coins to $G(\text{pk}, m)$. This is a countermeasure against multi-key attacks. In this paper we only model single-key security, so we omit pk from the hashes for brevity. The same also applies to the other transforms later in this paper, such as U^\perp .

3.2 Deterministic P : OW-CPA $\mathsf{P} \stackrel{\text{QROM}}{\Rightarrow} \text{IND-CCA } U^\perp(\mathsf{P}, \mathsf{F}, \mathsf{H})$

Our OW-CPA to IND-CCA conversion is in the style of [JZM19c]. However, that bound is based on the failure probability δ of a randomized encryption algorithm, whereas ours is based on the difficulty of finding a failure without access to the private key. This means our theorem applies to deterministic but imperfectly-correct algorithms, such as one of the three BIKE variants, BIKE-2 [ABB⁺19]. So instead we use injectivity and a game where the adversary tries to find ciphertexts which are valid but do not decrypt correctly.

Definition 2 (Valid ciphertext). *Let $\mathsf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ be a dPKE. Call a ciphertext c “valid” for a public key pk of P if there exists m such that $c = \text{Encr}(\text{pk}, m)$.*

We introduce a new failure-finding experiment⁶, to capture the probability that the adversary can find valid ciphertexts that cause a decryption failure.

Definition 3 (Finding Failing Ciphertext). *The find-failing-ciphertexts experiment (FFC) is shown in Figure 2. The FFC-advantage of an adversary \mathcal{A} is defined by*

$$\text{Adv}_{\mathsf{P}}^{\text{FFC}}(\mathcal{A}) := \Pr[\text{Expt}_{\mathsf{P}}^{\text{FFC}}(\mathcal{A}) \rightarrow 1].$$

The U^\perp transform converts a dPKE $\mathsf{P} = (\text{Keygen}_{\mathsf{P}}, \text{Encr}, \text{Decr})$ into a KEM $\mathsf{K} = (\text{Keygen}, \text{Encaps}, \text{Decaps})$ using a PRF $\mathsf{F} : \mathcal{K}_{\mathsf{F}} \times \mathcal{C} \rightarrow \mathcal{K}$ and a hash function $\mathsf{H} : \mathcal{M} \times \mathcal{C} \rightarrow \mathcal{K}$, modeled as random oracle [JZM19c, HHK17]. The PRF is used for implicit rejection, returning $\mathsf{F}(\text{prfk}, c)$ in case of an invalid ciphertext using a secret prfk . The U^\perp transform is defined in Figure 3. We also define variants $U_m^\perp, U^\perp, U_m^\perp$ of this transform, which make the following changes:

- On Encaps line 3 resp. Decaps line 7, the transformations U_m^\perp and U_m^\perp compute $\mathsf{H}(m)$ resp. $\mathsf{H}(m')$ instead of $\mathsf{H}(m, c)$ resp. $\mathsf{H}(m', c)$.
- On Decaps lines 4 and 6, the transformations U^\perp and U_m^\perp return \perp instead of $\mathsf{F}(\text{prfk}, c)$. These variants also don’t need a prfk as part of the private key.

The transforms U^\perp and U_m^\perp are said to use *explicit rejection* because they return an explicit failure symbol \perp . U^\perp and U_m^\perp are said to use *implicit rejection*.

⁶ It is a stretch to even call this an “experiment”, because it may not be possible to efficiently determine whether the adversary succeeded. In future work we hope to force the adversary to find failing *message*, but this version is simpler to integrate into our proof.

$\text{Expt}_P^{\text{FFC}}(\mathcal{A})$:

```

1   $H \xleftarrow{\$} \mathcal{H}$ 
2   $(pk, sk) \leftarrow \text{Keygen}()$ 
3   $L \leftarrow \mathcal{A}^H(pk)$ 
4  return  $[\exists m \in \mathcal{M}, c \in L : \text{Encr}(pk, m) = c \wedge \text{Decr}(sk, c) \neq m]$ 

```

Fig. 2. FFC experiment on a PKE P . The instantiation of H generalizes to any number of random oracles, including zero.

<u>Keygen():</u>	<u>Encaps(pk):</u>	<u>Decaps(sk, c):</u>
1 $(pk, sk_P) \leftarrow \text{Keygen}_P()$	1 $m \xleftarrow{\$} \mathcal{M}$	1 parse $sk = (sk_P, \text{prfk})$
2 $\text{prfk} \xleftarrow{\$} \mathcal{K}_F$	2 $c \leftarrow \text{Encr}(pk, m)$	2 $m' \leftarrow \text{Decr}(sk_P, c)$
3 $sk \leftarrow (sk_P, \text{prfk})$	3 $K \leftarrow H(m, c)$	3 if $m' = \perp$:
4 return (pk, sk)	4 return (K, c)	4 return $F(\text{prfk}, c)$
		5 else if $\text{Encr}(pk, m') \neq c$:
		6 return $F(\text{prfk}, c)$
		7 else: return $H(m', c)$

Fig. 3. Transform $U^\perp(P, F) := (\text{Keygen}, \text{Encaps}, \text{Decaps})$.

The next theorem states that breaking the IND-CCA security of $U^\perp(P, F, H)$ requires either breaking the OW-CPA security of P , causing a decapsulation failure, or breaking the PRF used for implicit rejection. In particular, we need P to be an ϵ -injective dPKE as in Definition 6.

Theorem 2. *Let $H : \mathcal{M} \times \mathcal{C} \rightarrow \mathcal{K}$ be a quantum-accessible random oracle, $F : \mathcal{K}_F \times \mathcal{C} \rightarrow \mathcal{K}$ be a PRF. Let P be an ϵ -injective dPKE which is independent of H . Let \mathcal{A} be an IND-CCA adversary against the KEM $U^\perp(P, F)$, and suppose that \mathcal{A} makes at most q_{dec} decryption queries. Then we can construct three adversaries running in about the same time and resources as \mathcal{A} :*

- an OW-CPA-adversary \mathcal{B}_1 against P
- a FFC-adversary \mathcal{B}_2 against P , returning a list of at most q_{dec} ciphertexts
- a PRF-adversary \mathcal{B}_3 against F

such that

$$\text{Adv}_{U^\perp(P)}^{\text{IND-CCA}}(\mathcal{A}) \leq 2\sqrt{\text{Adv}_P^{\text{OW-CPA}}(\mathcal{B}_1)} + \text{Adv}_P^{\text{FFC}}(\mathcal{B}_2) + 2 \cdot \text{Adv}_F^{\text{PRF}}(\mathcal{B}_3) + \epsilon.$$

In the common case that $F(\text{prfk}, c)$ is implemented as $H(\text{prfk}, c)$ it holds that if \mathcal{A} makes q queries at depth d , then

$$\text{Adv}_F^{\text{PRF}}(\mathcal{B}_3) \stackrel{\text{cor. 1}}{\leq} 2\sqrt{dq/|M|}.$$

Proof. Our proof is by a series of games. In some later games, we will define an outcome “draw” which is distinct from a win or loss. A draw counts as halfway between a win and a loss, as described by the adversary’s score w_i :

$$\begin{aligned} w_i &:= \Pr[\mathcal{A} \text{ wins: Game } i] + \frac{1}{2} \Pr[\text{Draw: Game } i] \\ &= \frac{1}{2} (\Pr[\mathcal{A} \text{ wins: Game } i] - \Pr[\mathcal{A} \text{ loses: Game } i]). \end{aligned}$$

Game 0 (IND-CCA). *This is the original IND-CCA game against the KEM $U^\perp(\mathsf{P}, \mathsf{F}, \mathsf{H})$, cf. Definition 12.*

Game 1 (PRF is random). *Game 1 is the same as Game 0, except the simulator replaces $\mathsf{F}(\text{prfk}, \cdot)$ with a random function $R \stackrel{\$}{\leftarrow} \mathcal{K}^{\mathcal{C}}$.*

We construct a PRF-adversary \mathcal{B}_3 (cf. Definition 10) which replaces its calls to $\mathsf{F}(\text{prfk}, \cdot)$ by calls to its oracle, runs \mathcal{A} , and outputs 1 if \mathcal{A} wins and 0 otherwise. Now, by construction $\Pr_{k \stackrel{\$}{\leftarrow} \mathcal{K}} [\mathcal{B}^{\mathsf{F}(k, \cdot)} = 1] = \Pr[\mathcal{A} \text{ wins: Game 0}]$ and $\Pr_{R \stackrel{\$}{\leftarrow} \mathcal{K}^{\mathcal{C}}} [\mathcal{A}^{R(\cdot)} = 1] = \Pr[\mathcal{A} \text{ wins: Game 1}]$. Hence,

$$|\Pr[\mathcal{A} \text{ wins: Game 1}] - \Pr[\mathcal{A} \text{ wins: Game 0}]| = \text{Adv}_F^{\text{PRF}}(\mathcal{A}).$$

Game 2 (Draw on fail or non-injective pk). *Let Fail be the event that one or more of \mathcal{A} ’s decapsulation queries $D(c)$ fails to decrypt, meaning that $c = \text{Encr}(\text{pk}, m)$ for some m , but $\text{Decr}(\text{sk}, c) \neq m$. Let NonInj be the event that $\text{Encr}(\text{pk}, \cdot)$ is not injective, and let $\text{Draw} := \text{Fail} \vee \text{NonInj}$. In Game 2 and onward, if Draw occurs then the game continues, but at the end it is a draw instead of the adversary winning or losing.*

Let $d_i := \Pr[\text{Draw : Game } i]$. Then $|w_2 - w_1| \leq \frac{1}{2}d_2$. Note that the event Draw is a well-defined classical event and does not depend on H , even though the simulator might not be able to determine efficiently whether it occurred.

Game 3 (Reprogram $H(m, c)$ to $R(c)$). *Game 3 is the same as Game 2, but the simulator reprograms $H(m, c)$ where $c = \text{Encr}(\text{pk}, m)$ to return $R(c)$.*

This produces the same win and draw probabilities as Game 2 as explained next. For each m , the value $H(m, \text{Encr}(\text{pk}, m))$ is changed to a uniformly, independently random value, except when the game is already a draw:

- It is uniformly random because R is uniformly random.
- It is independent of $H(m', c)$ for $m' \neq m$ because $\text{Encr}(\text{pk}, \cdot)$ is injective or else the game is a draw.
- H calls $R(c)$ only for valid ciphertexts $c = \text{Encr}(\text{pk}, m')$. On the other hand, the decapsulation oracle only calls $R(c')$ for rejected ciphertexts c' , i.e. ones where $c' \neq \text{Encr}(\text{pk}, \text{Decr}(\text{sk}, c'))$. If a valid ciphertext has been rejected and passed to R in this way, then Draw has occurred and the return value of R does not affect w_i or d_i .

Therefore $w_3 = w_2$ and $d_3 = d_2$.

Game 4 (Decapsulation oracle returns $R(c)$). *Game 4 is the same as Game 3, but the simulated decapsulation oracle simply returns $R(c)$ for all ciphertexts other than the challenge (for which it still returns \perp).*

In fact, the decapsulation oracle was already doing this in Game 3: The original decapsulation returns either $H(m, c)$ with $c = \text{Encr}(\text{pk}, m)$ or $F(\text{prfk}, c)$, but both of those have been reprogrammed to return $R(c)$. Therefore $w_4 = w_3$ and $d_4 = d_3$. As of this game, the simulator doesn't use the private key anymore.

Bound draw. We now want to upper bound the draw probability. Let \mathcal{B}_2 be the algorithm which, given a public key pk , simulates Game 4 for \mathcal{A} and outputs a list L of all of \mathcal{A} 's decapsulation queries. Then \mathcal{B}_2 is a FFC-adversary against P which runs in about the same time as \mathcal{A} and succeeds whenever a draw occurred during the game. Consequently,

$$d_2 = d_3 = d_4 \leq \text{Adv}_{\mathsf{P}}^{\text{FFC}}(\mathcal{B}_2) + \epsilon.$$

Game 5 (Change shared secret). *In Game 5, the shared secret is changed to a uniformly random value r . If $b = 1$, then for all m such that $\text{Encr}(\text{pk}, m) = c^*$, the oracle $H(m)$ is reprogrammed to return r . If $b = 0$, then H is not reprogrammed.*

If $\text{Encr}(\text{pk}, \cdot)$ is injective, then this is the same distribution as Game 4, and otherwise the game is a draw. Therefore $w_5 = w_4$.

It remains to bound \mathcal{A} 's advantage in Game 5. The simulation still runs in about the same time as \mathcal{A} . Furthermore, the $b = 0$ and $b = 1$ cases are now distinguished by a single return value from the H oracle. Hence, we can consider two oracles H and $H' := H[m^* \rightarrow r]$ as required by Lemma 5. Then Lemma 5, states that there is an algorithm \mathcal{B}_1 , running in about the same time as \mathcal{A} , such that for all H :

$$\begin{aligned} |\Pr[\text{Win} : b = 0] - \Pr[\text{Lose} : b = 1]| &= \left| \frac{\Pr[\mathcal{A} \rightarrow 0 \wedge \neg \text{Draw} : b = 0]}{-\Pr[\mathcal{A} \rightarrow 0 \wedge \neg \text{Draw} : b = 1]} \right| \\ &= \left| \frac{\Pr[\mathcal{A}^H \rightarrow 0 \wedge \neg \text{Draw}]}{-\Pr[\mathcal{A}^{H'} \rightarrow 0 \wedge \neg \text{Draw}]} \right| \\ &\leq 2\sqrt{\Pr[\mathcal{B}_1(\text{pk}, c) \rightarrow m^*]}. \end{aligned}$$

The same is true in expectation over H by Jensen's inequality:

$$\begin{aligned} &\mathbb{E} \left[2\sqrt{\Pr[\mathcal{B}_1(\text{pk}, c) \rightarrow m^*]} : H \xleftarrow{\$} \mathcal{K}^{(\mathcal{M} \times \mathcal{C})} \right] \\ &\leq 2\sqrt{\mathbb{E} \left[\Pr[\mathcal{B}_1(\text{pk}, c) \rightarrow m^*] : H \xleftarrow{\$} \mathcal{K}^{(\mathcal{M} \times \mathcal{C})} \right]} \\ &= 2\sqrt{\text{Adv}_{\mathsf{P}}^{\text{OW-CPA}}(\mathcal{B}_1)} \end{aligned}$$

so that

$$|\Pr[\text{Win} : b = 0] - \Pr[\text{Lose} : b = 1]| \leq 2\sqrt{\text{Adv}_{\mathcal{P}}^{\text{OW-CPA}}(\mathcal{B}_1)}.$$

Likewise, for the same adversary \mathcal{B}_1 ,

$$|\Pr[\text{Win} : b = 1] - \Pr[\text{Lose} : b = 0]| \leq 2\sqrt{\text{Adv}_{\mathcal{P}}^{\text{OW-CPA}}(\mathcal{B}_1)}.$$

Since b is either 0 or 1 each with probability $\frac{1}{2}$, we have by the triangle inequality:

$$|\Pr[\text{Win}] - \Pr[\text{Lose}]| \leq 2\sqrt{\text{Adv}_{\mathcal{P}}^{\text{OW-CPA}}(\mathcal{B}_1)}$$

so that $|w_5 - \frac{1}{2}| \leq \sqrt{\text{Adv}_{\mathcal{P}}^{\text{OW-CPA}}(\mathcal{B}_1)}$.

Summing up the differences in the previous games, we have

$$\left|w_0 - \frac{1}{2}\right| \leq \sqrt{\text{Adv}_{\mathcal{P}}^{\text{OW-CPA}}(\mathcal{B}_1)} + \frac{1}{2}\text{Adv}_{\mathcal{P}}^{\text{FFC}}(\mathcal{B}_2) + \frac{\epsilon}{2} + \text{Adv}_{\mathcal{F}}^{\text{PRF}}(\mathcal{B}_3)$$

and finally

$$\text{Adv}_{\mathcal{U} \neq (\mathcal{P})}^{\text{IND-CCA}}(\mathcal{A}) \leq 2\sqrt{\text{Adv}_{\mathcal{P}}^{\text{OW-CPA}}(\mathcal{B}_1)} + 2 \cdot \text{Adv}_{\mathcal{F}}^{\text{PRF}}(\mathcal{B}_3) + \text{Adv}_{\mathcal{P}}^{\text{FFC}}(\mathcal{B}_2) + \epsilon.$$

This completes the proof of Theorem 2. \square

Tightness This bound is essentially tight, since breaking the one-wayness of \mathcal{P} and finding decryption failures are both known to result in attacks. Breaking the PRF harms security if and only if implicit rejection is more secure than explicit rejection. For a correct \mathcal{P} the bound boils down to the first two terms of the sum. The square-root loss arises from OW being a weaker security notion than IND [MW18], i.e., harder to break, and recent results [JZM19b] suggest that the square-root loss might be unavoidable in the quantum setting.

3.3 Decryption failures

When the dPKE is constructed by derandomizing an rPKE, we can also bound the FFC advantage.

Lemma 6. *Let $\mathcal{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ be a δ -correct, ϵ -injective rPKE with messages in \mathcal{M} and randomness in \mathcal{R} . Let $G : \mathcal{M} \rightarrow \mathcal{R}$ be a random oracle, so that $T(\mathcal{P}, G) := (\text{Keygen}, \text{Encr}_1, \text{Decr})$ is a derandomized version of \mathcal{P} . Let \mathcal{A} be a FFC adversary against $T(\mathcal{P}, G)$ which makes at most q queries at depth d to G and returns a list of at most q_{dec} ciphertexts. Then*

$$\text{Adv}_{T(\mathcal{P}, G)}^{\text{FFC}}(\mathcal{A}) \leq ((4d + 1)\delta + \sqrt{3\epsilon}) \cdot (q + q_{\text{dec}}) + \epsilon.$$

Proof. See Appendix E.

Note that if ϵ is negligible, and if the adversary can recognize which ciphertexts will fail, then this is a Grover bound.

4 Explicit rejection and key confirmation

We now turn to systems with explicit rejection or key confirmation. The next theorem shows that the transform U^\perp (with explicit rejection) never yields KEMs that are more secure than KEMs constructed via U^\times (with implicit rejection).

Theorem 3 (Explicit \rightarrow implicit). *Let P be a dPKE. Let \mathcal{A} be an IND-CCA adversary against $U^\times(\mathsf{P}, \mathsf{F}, H)$. Then there is an IND-CCA adversary \mathcal{B} against $U^\perp(\mathsf{P}, H)$, running in about the same time and resources as \mathcal{B} , such that*

$$\text{Adv}_{U^\times(\mathsf{P}, \mathsf{F}, H)}^{\text{IND-CCA}}(\mathcal{A}) = \text{Adv}_{U^\perp(\mathsf{P}, H)}^{\text{IND-CCA}}(\mathcal{B}).$$

Proof. The only difference between $U^\perp(\mathsf{P}, H)$ and $U^\times(\mathsf{P}, \mathsf{F}, H)$ is that where the former would reject a ciphertext c by returning \perp , the latter instead returns $\mathsf{F}(\text{prfk}, c)$. So the adversary \mathcal{B} can simply choose a random PRF key prfk , run \mathcal{A} , and output \mathcal{A} 's result. \mathcal{B} forwards all of \mathcal{A} 's queries to its oracles and returns the responses with the only difference that in case the decapsulation oracle returns \perp , \mathcal{B} returns $\mathsf{F}(\text{prfk}, c)$. The algorithm \mathcal{B} perfectly simulates the IND-CCA game for $U^\times(\mathsf{P}, \mathsf{F}, H)$ and hence \mathcal{A} succeeds with the same success probability as in the original game. \square

On the other hand, explicit rejection is secure if key confirmation is used. Key confirmation refers to adding a hash of the message to the cipher text. Let τ be the number of bits desired for the key-confirmation tag. For a PKE $\mathsf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ define the transform $C(\mathsf{P}, H_t, \tau) := (\text{Keygen}, \text{Encr}_1, \text{Decr}_1)$ using a random oracle $H_t : \mathcal{M} \rightarrow \{0, 1\}^\tau$ as in Figure 4.

<u>Encr₁(pk, m):</u>	<u>Decr₁(sk, (c, t)):</u>
1 $c \leftarrow \text{Encr}(\text{pk}, m)$	1 $m' \leftarrow \text{Decr}(\text{sk}_{\mathsf{P}}, c)$
2 $t \leftarrow H_t(m)$	2 if $H_t(m') \neq t$:
3 return (c, t)	3 return \perp
	4 return m'

Fig. 4. Transform $C(\mathsf{P}, t) := (\text{Keygen}, \text{Encr}_1, \text{Decr}_1)$.

Theorem 4 (Implicit \rightarrow explicit with key confirmation). *Let P be an ϵ -injective dPKE. Consider the KEM $\mathsf{K}_1 := U_m^\perp(C(\mathsf{P}, H_t, \tau), H_s)$ obtained from P applying the C -transform with random oracle $H_t : \mathcal{M} \rightarrow \{0, 1\}^\tau$ and the U_m^\perp -transform with independent random oracle $H_s : \mathcal{M} \rightarrow \{0, 1\}^s$. Let $\mathsf{K}_2 := U_m^\times(\mathsf{P}, \mathsf{F}, H)$ be the KEM obtained from P applying the U_m^\times -transform with random oracle $H : \mathcal{M} \rightarrow \{0, 1\}^{s+\tau}$.*

If \mathcal{A} is an IND-CCA-adversary against K_1 which makes q_{dec} decapsulation queries, then it is also an IND-CCA-adversary against P_2 and there is a PRF-adversary \mathcal{B} against F which uses about the same time and resources as \mathcal{A} , such

that:

$$\text{Adv}_{\mathcal{K}_1}^{\text{IND-CCA}}(\mathcal{A}) \leq 2 \cdot \text{Adv}_{\mathcal{K}_2}^{\text{IND-CCA}}(\mathcal{A}) + \frac{q_{\text{dec}}}{2^{\tau-1}} + 2 \cdot \text{Adv}_{\mathbb{F}}^{\text{PRF}}(\mathcal{B}) + 2\epsilon.$$

Proof. The gist of the proof is that the tag and the shared secret are both computed by hashing m . This means we can emulate \mathcal{K}_1 , where the ζ -bit shared secret and τ -bit tag are computed by splitting the $(\zeta + \tau)$ -bit shared secret of \mathcal{K}_2 . In the case that \mathcal{K}_2 implicitly rejects a ciphertext, then it returns a pseudorandom key. The tag parsed from this key will not match the adversary's tag except with probability $2^{-\tau}$, so the emulated \mathcal{K}_1 will almost always reject as well.

We defer the details to Appendix F.

Finally, we can show that hashing m is equivalent to hashing (m, c) in the next theorem.

Theorem 5 ($U_m \leftrightarrow U$). *Let \mathcal{P} be a $d\text{PKE}$. Let $\mathcal{K}_1 = U^\perp(\mathcal{P}, H_1)$ and $\mathcal{K}_2 = U_m^\perp(\mathcal{P}, H_2)$. Then \mathcal{K}_1 is IND-CCA secure if and only if \mathcal{K}_2 is IND-CCA secure. In other words, if there is an adversary \mathcal{A} against one, then there is an adversary \mathcal{B} against the other, running in about the same time and with the same advantage.*

The same is true for U^χ and U_m^χ .

Proof. This is a simple indistinguishability argument. In both the encapsulation and decapsulation functions, the IND-CCA experiment against \mathcal{K}_1 only calls $H_1(m, c)$ when $c = \text{Encr}(\text{pk}, m)$. So to simulate the \mathcal{K}_1 -experiment playing in an IND-CCA experiment against \mathcal{K}_2 (with oracle $H_2 : \mathcal{M} \rightarrow \mathcal{K}$), sample fresh random oracle $H \xleftarrow{\$} \mathcal{K}^{(\mathcal{M}, c)}$ and set

$$H_1(m, c) := \begin{cases} H_2(m) & \text{if } c = \text{Encr}(\text{pk}, m), \\ H(m, c) & \text{otherwise.} \end{cases}$$

This exactly simulates the IND-CCA experiment against \mathcal{K}_1 . In the other direction, to simulate the IND-CCA experiment against \mathcal{K}_2 it suffices to redirect $H_2(m)$ to $H_1(m, \text{Encr}(\text{pk}, m))$.

The same technique works for U^χ and U_m^χ . It also works for security notions other than IND-CCA, such as OW-CCA, OW-qPVCA, etc. \square

5 Acknowledgements

Part of this work was done while the authors were participating in the 2019 Oxford Post-Quantum Cryptography Workshop.

Special thanks to Daniel J. Bernstein, Edward Eaton and Mark Zhandry for helpful discussions.

References

- ABB⁺19. Nicolas Aragon, Paulo S. L. M. Barreto, Slim Bettaiieb, Loïc Bidoux, Olivier Blazy, Jean-Christophe Deneuville, Philippe Gaborit, Shay Gueron, Tim

- Güneysu, Carlos Aguilar Melchor, Rafael Misoczki, Edoardo Persichetti, Nicolas Sendrier, Jean-Pierre Tillich, Valentin Vasseur, and Gilles Zémor. BIKE: Bit flipping key encapsulation, 2019. <https://bikesuite.org>.
- AHU18. Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. Cryptology ePrint Archive, Report 2018/904, 2018. <https://eprint.iacr.org/2018/904>.
- BDF⁺11. Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer, Heidelberg, December 2011. doi:10.1007/978-3-642-25385-0_3.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993. doi:10.1145/168588.168596.
- FO99. Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 537–554. Springer, Heidelberg, August 1999. doi:10.1007/3-540-48405-1_34.
- HHK17. Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 341–371. Springer, Heidelberg, November 2017. doi:10.1007/978-3-319-70500-2_12.
- HNP⁺03. Nick Howgrave-Graham, Phong Q. Nguyen, David Pointcheval, John Proos, Joseph H. Silverman, Ari Singer, and William Whyte. The impact of decryption failures on the security of NTRU encryption. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 226–246. Springer, Heidelberg, August 2003. doi:10.1007/978-3-540-45146-4_14.
- JZC⁺17. Haodong Jiang, Zhenfeng Zhang, Long Chen, Hong Wang, and Zhi Ma. Post-quantum IND-CCA-secure KEM without additional hash. Cryptology ePrint Archive, Report 2017/1096, 2017. <https://eprint.iacr.org/2017/1096>.
- JZM19a. Haodong Jiang, Zhenfeng Zhang, and Zhi Ma. Key encapsulation mechanism with explicit rejection in the quantum random oracle model. Cryptology ePrint Archive, Report 2019/052, 2019. <https://eprint.iacr.org/2019/052>.
- JZM19b. Haodong Jiang, Zhenfeng Zhang, and Zhi Ma. On the non-tightness of measurement-based reductions for key encapsulation mechanism in the quantum random oracle model. Cryptology ePrint Archive, Report 2019/494, 2019. <https://eprint.iacr.org/2019/494>.
- JZM19c. Haodong Jiang, Zhenfeng Zhang, and Zhi Ma. Tighter security proofs for generic key encapsulation mechanism in the quantum random oracle model. Cryptology ePrint Archive, Report 2019/134, 2019. <https://eprint.iacr.org/2019/134>.
- MW18. Daniele Micciancio and Michael Walter. On the bit security of cryptographic primitives. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 3–28. Springer, Heidelberg, April / May 2018. doi:10.1007/978-3-319-78381-9_1.
- NC00. Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

- SXY18. Tsunekazu Saito, Keita Xagawa, and Takashi Yamakawa. Tightly-secure key-encapsulation mechanism in the quantum random oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 520–551. Springer, Heidelberg, April / May 2018. doi:10.1007/978-3-319-78372-7_17.
- TU16. Ehsan Ebrahimi Targhi and Dominique Unruh. Post-quantum security of the Fujisaki-Okamoto and OAEP transforms. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 192–216. Springer, Heidelberg, October / November 2016. doi:10.1007/978-3-662-53644-5_8.
- Unr15. Dominique Unruh. Revocable quantum timed-release encryption. *J. ACM*, 62(6):49:1–49:76, December 2015. URL: <http://doi.acm.org/10.1145/2817206>, doi:10.1145/2817206.
- XY18. Keita Xagawa and Takashi Yamakawa. (tightly) qcca-secure key-encapsulation mechanism in the quantum random oracle model. Cryptology ePrint Archive, Report 2018/838, 2018. <https://eprint.iacr.org/2018/838>.
- Zha18. Mark Zhandry. How to record quantum queries, and applications to quantum indifferenciability. Cryptology ePrint Archive, Report 2018/276, 2018. <https://eprint.iacr.org/2018/276>.

A Security Notions and Definitions

In this section we recall the definition of KEMs and PKEs. Additionally, we recall the respective security notions that are needed in this paper.




Fig. 5. Definition of the classical and quantum random oracle

Definition 4 (Public-Key Encryption Schemes). A randomized *public-key encryption scheme* (*rPKE*) is defined over a finite message space \mathcal{M} , a ciphertext space \mathcal{C} , a secret key space \mathcal{SK} and a public key space \mathcal{PK} . It consists of a triple of algorithms $\mathsf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ defined as follows.

- $\text{Keygen}() \rightarrow (\text{pk}, \text{sk})$ is a randomized algorithm that returns a secret key $\text{sk} \in \mathcal{SK}$ and a public key $\text{pk} \in \mathcal{PK}$.
- $\text{Encr}(\text{pk}, \text{m}) \rightarrow c$ is a randomized algorithm that takes as input a public key pk and a message $\text{m} \in \mathcal{M}$, and outputs a ciphertext $c \in \mathcal{C}$.

- $\text{Decr}(\text{sk}, c) \rightarrow \{m', \perp\}$ is a deterministic algorithm that takes as input a secret key $\text{sk} \in \mathcal{SK}$ and a ciphertext $c \in \mathcal{C}$ and returns either a message $m' \in \mathcal{M}$ or a failure symbol $\perp \notin \mathcal{M}$.

A deterministic public-key encryption scheme (dPKE) is defined the same way, except that Encr is a deterministic algorithm.

Definition 5 (Correctness and failure probability of PKEs). A PKE $\mathsf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ is δ -correct if

$$\mathbb{E} \left[\max_{m \in \mathcal{M}} \Pr[\text{Decr}(\text{sk}, \text{Encr}(\text{pk}, m)) \neq m] : (\text{pk}, \text{sk}) \leftarrow \text{Keygen}() \right] \leq \delta.$$

We call δ the decryption failure probability of P . We say P is correct if $\delta = 0$.

Note that this definition works for a deterministic or randomized PKE, but for a deterministic PKE the term $\max_{m \in \mathcal{M}} \Pr[\text{Decr}(\text{sk}, \text{Encr}(\text{pk}, m)) \neq m]$ is either 0 or 1 for each keypair.

Definition 6 (Injectivity of PKEs). A dPKE $\mathsf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ is ϵ -injective if

$$\Pr \left[\text{Encr}(\text{pk}, m) \text{ is not injective} : (\text{pk}, \text{sk}) \leftarrow \text{Keygen}(), H \leftarrow^{\$} \mathcal{H} \right] \leq \epsilon.$$

We say P is injective if $\epsilon = 0$.

Definition 7 (OW-CPA Advantage). Let $\mathsf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ be a dPKE or rPKE. The one-way under chosen-plaintext attacks (OW-CPA) experiment is shown in Figure 6. The OW-CPA-advantage of an adversary \mathcal{A} is defined as

$$\text{Adv}_{\mathsf{P}}^{\text{OW-CPA}}(\mathcal{A}) := \Pr[\text{Expt}_{\mathsf{P}}^{\text{OW-CPA}}(\mathcal{A}) \rightarrow 1].$$

Note that some papers, e.g., [JZM19c], define OW-CPA-advantage this way, and some, e.g., [HHK17], instead use the looser condition that $\text{Encr}(\text{pk}, m') = c^*$, particularly if Encr is deterministic. We use the definition in Figure 6 because it is more convenient for our proofs of Theorems 1 and 2.

Definition 8 (IND-CPA Advantage). Let $\mathsf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ be an rPKE. The indistinguishability under chosen-plaintext attacks (IND-CPA) experiment is shown in Figure 6. The IND-CPA-advantage of an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is defined as

$$\text{Adv}_{\mathsf{P}}^{\text{IND-CPA}}(\mathcal{A}) := 2 \left| \Pr[\text{Expt}_{\mathsf{P}}^{\text{IND-CPA}}(\mathcal{A}) \rightarrow 1] - \frac{1}{2} \right|.$$

Note that IND-CPA is unachievable for dPKEs, because \mathcal{A} can just test which message encrypts to c^* .

A weakening of IND-CPA is IND-KPA where the challenge messages are chosen by the experiment.

Definition 9 (IND-KPA Advantage). Let $P = (\text{Keygen}, \text{Encr}, \text{Decr})$ be an rPKE. The indistinguishability under known-plaintext attack (IND-KPA) experiment is shown in Figure 6. The IND-KPA-advantage of an adversary \mathcal{A} is defined as

$$\text{Adv}_P^{\text{IND-KPA}}(\mathcal{A}) := 2 \left| \Pr[\text{Expt}_P^{\text{IND-KPA}}(\mathcal{A}) \rightarrow 1] - \frac{1}{2} \right|.$$

Clearly $\text{IND-CPA} \Rightarrow \text{IND-KPA}$ as any IND-KPA adversary can be used to break IND-CPA using it as \mathcal{A}_2 and simulating \mathcal{A}_1 by just sampling random messages.

$\text{Expt}_P^{\text{OW-CPA}}(\mathcal{A})$:	$\text{Expt}_P^{\text{IND-CPA}}(\mathcal{A})$:
1 $H \xleftarrow{\$} \mathcal{H}$	1 $H \xleftarrow{\$} \mathcal{H}$
2 $(\text{pk}, \text{sk}) \leftarrow \text{Keygen}()$	2 $(\text{pk}, \text{sk}) \leftarrow \text{Keygen}()$
3 $m^* \xleftarrow{\$} \mathcal{M}$	3 $(st, m_0, m_1) \leftarrow \mathcal{A}_1^H(\text{pk})$
4 $c^* \leftarrow \text{Encr}(\text{pk}, m^*)$	4 $b \xleftarrow{\$} \{0, 1\}$
5 $m' \leftarrow \mathcal{A}^H(\text{pk}, c^*)$	5 $c^* \leftarrow \text{Encr}(\text{pk}, m_b^*)$
6 return $[m^* = m']$	6 $b' \leftarrow \mathcal{A}_2^H(\text{pk}, m_0, m_1, c^*, st)$
	7 return $[b = b']$
$\text{Expt}_P^{\text{IND-KPA}}(\mathcal{A})$:	
1 $H \xleftarrow{\$} \mathcal{H}$	
2 $(\text{pk}, \text{sk}) \leftarrow \text{Keygen}()$	
3 $m_0, m_1 \leftarrow \mathcal{M}$	
4 $b \xleftarrow{\$} \{0, 1\}$	
5 $c^* \leftarrow \text{Encr}(\text{pk}, m_b^*)$	
6 $b' \leftarrow \mathcal{A}^H(\text{pk}, m_0, m_1, c^*)$	
7 return $[b = b']$	

Fig. 6. OW-CPA, IND-CPA, and IND-KPA of a PKE P . The instantiation of H generalizes to any number of random oracles, including zero.

Definition 10 (PRF Advantage). Let $F : \mathcal{K}_F \times X \rightarrow Y$ be a pseudorandom function (PRF). We define the PRF-advantage of an adversary \mathcal{A} as

$$\text{Adv}_F^{\text{PRF}}(\mathcal{A}) = \left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}^{F(k, \cdot)} = 1 \right] - \Pr_{R \leftarrow \mathcal{S}^{Y \times X}} \left[\mathcal{A}^{R(\cdot)} = 1 \right] \right|.$$

Definition 11 (Key Encapsulation Mechanism). A KEM K defined over the message space \mathcal{M} , the public key space \mathcal{PK} , the secret key space \mathcal{SK} , and the key space \mathcal{K} , is a triple of algorithms $K = (\text{Keygen}, \text{Encaps}, \text{Decaps})$ defined as follows.

- $\text{Keygen}() \rightarrow (\text{pk}, \text{sk})$ is a randomized algorithm that returns a public key $\text{pk} \in \mathcal{PK}$ and a secret key $\text{sk} \in \mathcal{SK}$.

- $\text{Encaps}(\text{pk}) \rightarrow (c, \kappa)$ is a randomized algorithm that takes as input a public key pk and outputs a ciphertext c as well as a key $\kappa \in K$.
- $\text{Decaps}(\text{sk}, c) \rightarrow \kappa$ or \perp is a deterministic algorithm that takes as input a secret key $\text{sk} \in \mathcal{SK}$ and a ciphertext c and returns a key $\kappa \in K$ or a failure symbol $\perp \notin K$.

As before, we use \mathcal{H} to denote the space of functions from which the random hash function is randomly sampled if a proof for K is being given in the ROM.

Definition 12 (IND-CCA Advantage). Let K be a KEM. The security experiment $\text{Expt}_{\mathsf{K}}^{\text{IND-CCA}}(\mathcal{A})$ is defined in Figure 7 for an adversary \mathcal{A} against K , given access to a (quantum-accessible) random oracle H and a classical decapsulation oracle D .

We define the advantage of a classical (resp., quantum) adversary \mathcal{A} against a KEM K in the classical (resp., quantum-accessible) random oracle model as

$$\text{Adv}_{\mathsf{K}}^{\text{IND-CCA}}(\mathcal{A}) = \left| \Pr \left[\text{Expt}_{\mathsf{K}}^{\text{IND-CCA}}(\mathcal{A}) = 1 \right] - \frac{1}{2} \right|.$$

$\text{Expt}_{\mathsf{K}}^{\text{IND-CCA}}(\mathcal{A})$:	Classical decapsulation oracle $D(c)$:
1 $H \xleftarrow{\$} \mathcal{H}$	1 if $c = c^*$: return \perp
2 $(\text{pk}, \text{sk}) \leftarrow \text{Keygen}()$	2 return $\text{Decaps}(\text{sk}, c)$
3 $(c^*, k_0^*) \leftarrow \text{Encaps}(\text{pk})$	
4 $k_1^* \xleftarrow{\$} K$	
5 $b \xleftarrow{\$} \{0, 1\}$	
6 $b' \leftarrow \mathcal{A}^{H, D}(\text{pk}, c^*, k_b^*)$	
7 return $[b = b']$	

Fig. 7. IND-CCA security experiment in the (Q)ROM against an adversary \mathcal{A}

B Proof of Lemma 5

Lemma 5 (Double-sided O2H). Let $G, H : X \rightarrow Y$ be random functions, let z be a random value, and let $S \subset X$ be a random set such that $\forall x \notin S, G(x) = H(x)$. (G, H, S, z) may have arbitrary joint distribution. Let $f : X \rightarrow W \subseteq \{0, 1\}^n$ be a function such that

$$\{f(x) : x \in S\} = \{w^*\}$$

is a single element. Let \mathcal{A}^H be a quantum oracle algorithm. Let Ev be an arbitrary classical event. We will define another quantum oracle algorithm $\mathcal{B}^{G, H}(z)$ which returns either w^* or a failure symbol \perp . \mathcal{B} runs in about the same amount of

time as \mathcal{A} , but when \mathcal{A} queries H , \mathcal{B} queries both G and H , and also runs f twice. Let

$$P_{\text{left}} := \Pr[\text{Ev} : \mathcal{A}^H(z)], \quad P_{\text{right}} := \Pr[\text{Ev} : \mathcal{A}^G(z)], \quad P_{\text{extract}} := \Pr[\mathcal{B}^{G,H}(z) \rightarrow w^*].$$

Then

$$|P_{\text{left}} - P_{\text{right}}| \leq 2\sqrt{P_{\text{extract}}} \quad \text{and} \quad \left| \sqrt{P_{\text{left}}} - \sqrt{P_{\text{right}}} \right| \leq 2\sqrt{P_{\text{extract}}}.$$

Proof. The outline of our proof is to use the compressed oracle framework from [Zha18] to instantiate a new random oracle $B : W \rightarrow \{0, 1\}$. On query x , the simulator returns $H(x)$ if $B(f(x)) = 0$, or $G(x)$ if $B(f(x)) = 1$. The only value of $B(z)$ that affects the result is $b := B(w^*)$, so at the end of the computation the compressed oracle table for B must be either the empty table or $\{w^* \rightarrow b\}$. Our proof quantizes and simplifies this outline.

To begin, suppose that G and H are fixed and \mathcal{A} is unitary. Consider an algorithm $\mathcal{B}_0^{H,G}$ which runs \mathcal{A}^H and \mathcal{A}^G in superposition, with an additional bit b signifying which oracle is being used. Then if \mathcal{A}^G behaves differently from \mathcal{A}^H , the state of \mathcal{A} will become entangled with b . We will use $|b\rangle = |+\rangle := (|0\rangle + |1\rangle)/\sqrt{2}$ to signify that \mathcal{A} is using H , and $|b\rangle = |-\rangle := (|0\rangle - |1\rangle)/\sqrt{2}$ to signify that \mathcal{A} is using G . That is:

$$\mathcal{B}_0^{H,G} := \frac{\mathcal{A}^H \otimes |+\rangle + \mathcal{A}^G \otimes |-\rangle}{\sqrt{2}}.$$

This \mathcal{B}_0 can be implemented as \mathcal{A} with only the oracle queries changed. To do this, let b start in the state $(|+\rangle + |-\rangle)/\sqrt{2} = |0\rangle$. When \mathcal{A} queries the oracle, \mathcal{B}_0 implements the following map on $|x, y, b\rangle$:

$$\begin{aligned} U(|x, y, +\rangle) &:= |x, y \oplus H(x), +\rangle, \\ U(|x, y, -\rangle) &:= |x, y \oplus G(x), -\rangle. \end{aligned}$$

This is the same as a conditional evaluation map which queries H if $b = 0$ and G if $b = 1$, with a Hadamard transform before and after.

Let ψ_H resp. ψ_G be the final states of \mathcal{A}^H resp. \mathcal{A}^G . The final state of $\mathcal{B}_0^{H,G}$ is

$$\frac{\psi_H \otimes |+\rangle + \psi_G \otimes |-\rangle}{\sqrt{2}} = \frac{1}{2} \cdot \left((|\psi_H\rangle + |\psi_G\rangle) \otimes |0\rangle + (|\psi_H\rangle - |\psi_G\rangle) \otimes |1\rangle \right).$$

Suppose we measure b in the computational basis. This commutes with the final measurement of \mathcal{A} 's state. Then, we will measure 1 with probability $\epsilon := \|\psi_H - \psi_G\|^2/4$, and hence,

$$\|\psi_H - \psi_G\| = 2\sqrt{\epsilon}.$$

This $2\sqrt{\epsilon}$ is the claimed probability bound, but we still need a way to extract w^* . The full algorithm $\mathcal{B}^{H,G}$ is the same as \mathcal{B}_0 , but with a different final measurement and another auxiliary register $w \in \{0, 1\}^n$ (i.e. a register that can represent elements of w). The w register is initialized to 0.

We will ensure that except during queries, (b, w) will always be in the state $(0, 0)$ or $(1, w^*)$. More formally, let T_w operate on the b and w registers by $T_w(|b, w\rangle) := |b, w \oplus (b \cdot w^*)\rangle$. Therefore T_w swaps $(1, 0)$ with $(1, w^*)$. We will ensure that if at some step in the computation the state of $\mathcal{B}_0^{H,G}$ is ψ , then during the same step the state of $\mathcal{B}^{H,G}$ is $T_w(\psi \otimes |0\rangle)$.

Since $\mathcal{B}_0^{H,G}$ replaces the oracle queries with U , $\mathcal{B}^{H,G}$ should replace them with $U_w := T_w \circ U \circ T_w^\dagger$. (This gives the desired result because T_w commutes with all the steps of \mathcal{A} except for the oracle queries.) To do this, let

$$T_f(|x, y, b, w\rangle) := |x, y, b, w \oplus (b \cdot f(x))\rangle.$$

Then $\mathcal{B}^{H,G}$ replaces \mathcal{A} 's oracle queries with

$$U_f := T_f \circ U \circ T_f^\dagger.$$

In fact, $U_f = U_w$. On the subspace where $x \in S$, we have $f(x) = w^*$ by assumption. Therefore $T_f = T_w$ and $U_f = U_w$. On the orthogonal subspace where $x \notin S$, we have $G(x) = H(x)$, so the operation U does not depend on b or w . Therefore on that subspace, U commutes with T_f and T_w , so that $U_f = U = U_w$. In sum, $U_f = U_w$ is an efficient implementation of the oracle by $\mathcal{B}^{H,G}$.

When \mathcal{A} completes, $\mathcal{B}^{H,G}$ measures (b, w) in the computational basis. With probability ϵ it measures $(1, w^*)$, in which case it outputs w^* . Otherwise it measures $(0, 0)$, in which case it outputs \perp .

By [AHU18] Lemmas 3 and 4,

$$|\Pr[\mathcal{A}^H \rightarrow 1] - \Pr[\mathcal{A}^G \rightarrow 1]| \leq \| |\psi_0\rangle - |\psi_1\rangle \| \leq 2\sqrt{\Pr[\mathcal{B}^{H,G} \rightarrow w^*]}$$

and likewise

$$\left| \sqrt{\Pr[\mathcal{A}^H \rightarrow 1]} - \sqrt{\Pr[\mathcal{A}^G \rightarrow 1]} \right| \leq \| |\psi_0\rangle - |\psi_1\rangle \| \leq 2\sqrt{\Pr[\mathcal{B}^{H,G} \rightarrow w^*]}.$$

This completes the proof for unitary adversaries \mathcal{A} with a fixed H and G .

For non-unitary adversaries and for random distributions of H, G , we instead end in a mixture of states Ψ_0 resp. Ψ_1 , for which Euclidean distance is not appropriate but the Bures distance [NC00] is. By monotonicity and joint concavity of fidelity (exactly as in [AHU18] Lemma 6 and 9), the same bound holds for the Bures distance:

$$|\Pr[\mathcal{A}^H \rightarrow 1] - \Pr[\mathcal{A}^G \rightarrow 1]| \leq B(\Psi_0, \Psi_1) \leq 2\sqrt{\Pr[\mathcal{B}^{H,G} \rightarrow w^*]}$$

and likewise

$$\left| \sqrt{\Pr[\mathcal{A}^H \rightarrow 1]} - \sqrt{\Pr[\mathcal{A}^G \rightarrow 1]} \right| \leq B(\Psi_0, \Psi_1) \leq 2\sqrt{\Pr[\mathcal{B}^{H,G} \rightarrow w^*]}.$$

This completes the proof in the general case. \square

C Proof of Theorem 1

Theorem 1. *Let P be an $r\text{PKE}$ with messages in \mathcal{M} and random coins in \mathcal{R} . Let $G : \mathcal{M} \rightarrow \mathcal{R}$ be a quantum-accessible random oracle. Let \mathcal{A} be an OW-CPA adversary against $\mathsf{P}' := T(\mathsf{P}, G)$. Suppose that \mathcal{A} queries G at most q times with depth at most d .*

Then we can construct an IND-CPA adversary \mathcal{B} against P , running in about the same time and resources as \mathcal{A} , such that

$$\text{Adv}_{\mathsf{P}'}^{\text{OW-CPA}}(\mathcal{A}) \leq (d+2) \cdot \left(\text{Adv}_{\mathsf{P}}^{\text{IND-CPA}}(\mathcal{B}) + \frac{8(q+1)}{|\mathcal{M}|} \right).$$

Proof. Let \mathcal{A}_1 be the same as \mathcal{A} , except that at the end after choosing an output m , it computes and discards $G(m)$. Therefore it makes at most $q+1$ queries at depth at most $d+1$. This is a formality so that returning the correct m will count as a Find (see Definition 1) later in the proof. Clearly the two algorithms \mathcal{A} and \mathcal{A}_1 have the same OW-CPA-advantage against P' .

We actually show a slightly stronger result, constructing an IND-KPA adversary \mathcal{B} . The IND-KPA adversary \mathcal{B} (cf. Definition 9) is given the tuple

$$(\text{pk}, m_0, m_1, c) \text{ where } c = \text{Encr}(\text{pk}, m_b; r).$$

It wants to determine whether $b = 0$ or $b = 1$. The algorithm \mathcal{B} creates a fresh random oracle G and runs

$$\mathcal{A}_1^{G \setminus \{m_0, m_1\}}(\text{pk}, c).$$

Suppose Find occurs, i.e., a query $x \in \{m_0, m_1\}$ was asked by \mathcal{A} to its oracle G . Then \mathcal{B} measures whether the query x was m_0 or m_1 , and returns the corresponding b . If Find does not occur, or if Find occurs but both m_0 and m_1 were queried, then \mathcal{B} guesses b at random.

Let G' be the oracle such that $G'(m_b) = r$ gives the encryption coins used to encrypt m_b , but $G'(m) = G(m)$ for all other messages m . G' is unknown to \mathcal{B} , but we can still analyze \mathcal{A} 's behavior when run with G' instead of G .

By construction, $\mathcal{A}_1^{G' \setminus \{m_0, m_1\}}$ cannot return m_b without causing Find. Hence,

$$\begin{aligned} \sqrt{\text{Adv}_{\mathsf{P}'}^{\text{OW-CPA}}(\mathcal{A})} &= \sqrt{\Pr[\mathcal{A}^{G'} \rightarrow m_b]} \\ &= \left| \sqrt{\Pr[\mathcal{A}_1^{G'} \rightarrow m_b]} \right| - \underbrace{\sqrt{\Pr[\mathcal{A}_1^{G' \setminus \{m_0, m_1\}} \rightarrow m_b \wedge \neg \text{Find}]}}_{=0} \\ &\stackrel{\text{Lem. 3}}{\leq} \sqrt{(d+2) \cdot \Pr[\text{Find} : \mathcal{A}_1^{G' \setminus \{m_0, m_1\}}]}. \end{aligned}$$

Squaring both sides,

$$\begin{aligned} \text{Adv}_{\mathsf{P}'}^{\text{OW-CPA}}(\mathcal{A}) &\leq (d+2) \cdot \Pr[\text{Find} : \mathcal{A}_1^{G' \setminus \{m_0, m_1\}}] \\ &\stackrel{\text{Lem. 2}}{=} (d+2) \cdot \Pr[\text{Find} : \mathcal{A}_1^{G \setminus \{m_0, m_1\}}] \\ &= (d+2) \cdot \Pr[\text{Find} : \mathcal{B}]. \end{aligned}$$

Now decompose Find as $\text{Find}_b \vee \text{Find}_{-b}$, where the former event means that Find occurs and m_b is measured, and latter means that Find occurs and m_{-b} is measured. They can both occur if \mathcal{A} makes multiple queries simultaneously, but $\text{Adv}_{\mathcal{P}}^{\text{IND-KPA}}(\mathcal{B}) = |\Pr[\text{Find}_b] - \Pr[\text{Find}_{-b}]|$ regardless.

Moreover, since \mathcal{B} measures m whenever Find occurs, we can view $G \setminus \{m_0, m_1\}$ as $G'' \setminus \{m_{-b}\} := (G \setminus \{m_b\}) \setminus \{m_{-b}\}$. Since \mathcal{A} has no information about m_{-b} except from puncturing, it holds for any m that

$$\Pr \left[m \in \{m_{-b}\} : \mathcal{A}^{G''} \right] = 1/|\mathcal{M}| =: \epsilon.$$

By (the second statement in) Lemma 4, we have

$$\Pr[\text{Find}_{-b} : \mathcal{B}] \leq 4(q+1)\epsilon = \frac{4(q+1)}{|\mathcal{M}|}.$$

Hence,

$$\begin{aligned} \text{Adv}_{\mathcal{P}'}^{\text{IND-KPA}}(\mathcal{B}) &= |\Pr[\text{Find}_b : \mathcal{B}] - \Pr[\text{Find}_{-b} : \mathcal{B}]| \\ &\geq \Pr[\text{Find} : \mathcal{B}] - 2\Pr[\text{Find}_{-b} : \mathcal{B}] \\ &\geq \Pr[\text{Find} : \mathcal{B}] - 8(q+1)/|\mathcal{M}|. \end{aligned}$$

Taking into account that $\text{Adv}_{\mathcal{P}}^{\text{IND-KPA}}(\mathcal{B}) \leq \text{Adv}_{\mathcal{P}}^{\text{IND-CPA}}(\mathcal{B})$ and combining these results gives

$$\text{Adv}_{\mathcal{P}'}^{\text{OW-CPA}}(\mathcal{A}) \leq (d+2) \cdot \left(\text{Adv}_{\mathcal{P}}^{\text{IND-CPA}}(\mathcal{B}) + \frac{8(q+1)}{|\mathcal{M}|} \right)$$

as claimed. \square

IND vs. OW. Our \mathcal{B} is a distinguishing adversary, not a one-way adversary. The reason is that \mathcal{A} can check whether a given m is the challenge message, but if \mathcal{P} is semantically secure then \mathcal{B} cannot check this. Instead \mathcal{B} would have to pick a random query to measure, which still works using Lemma 1, but with an additional factor of q tightness loss. That is, the one-way problem is potentially harder for a randomized encryption scheme than for a deterministic one. The authors discussed using a new “one-way with confirmation oracle” security game to more tightly capture the OW vs. IND tradeoff, but decided that it is simpler to just reduce to IND-CPA.

We also note that ordinarily distinguishing adversaries are much harder to amplify than one-way adversaries, but \mathcal{B} is constructed to either output with relative certainty if Find , or to fail and guess at random. This means that its advantage will still be high in the Micciancio-Walter notion of cryptographic advantage [MW18]. It is likely that the $8(q+1)/|\mathcal{M}|$ could be reduced to a $4(q+1)/|\mathcal{M}|$ without this requirement.

D Why encryption is usually injective for LWE

Here we outline why we expect $\text{Encr}(\text{pk}, \cdot)$ to be an injective function for the overwhelming majority of public keys pk in a PKE based on Learning with Errors (LWE). Consider a typical LWE PKE, where the public key has the form $(A, S = sA + e)$ where s and e are small of dimension n , and ciphertexts have the form $(As' + e', \lceil Xs' + e' + \text{encode}(m) \rceil)$. Encryption will fail to be injective for some G if there are $(s'_0, e'_0) = G(m_0)$ and $(s'_1, e'_1) = G(m_1)$ such that

$$As'_0 + e'_0 = As'_1 + e'_1 \quad \text{and} \quad \lceil Xs'_0 + e'_0 + \text{encode}(m_0) \rceil = \lceil Xs'_1 + e'_1 + \text{encode}(m_1) \rceil.$$

For correctness, the rounded component is always larger than the message space, and is generally larger than $|M|^2$. The unrounded component has size at least q^n which is larger still. The function $(s'_0, e'_0) \rightarrow As'_0 + e'_0$ is almost a universal hash unless s'_0 has large nullity, which is highly unlikely for any secure PKE. So the probability of collision with fewer than $|M|^2$ message pairs is not much bigger than q^{-n} , which is negligible.

E Proof of Lemma 6

To prove Lemma 6, we first show a result about Bernoulli variables.

Lemma 7. *Let $\{e_i : 1 \leq i \leq n\}$ be a collection of n independent Bernoulli variables. Let $\delta := \max \Pr[e_i]$, and for each integer j let $p_j := \Pr[\sum_{i=1}^n e_i = j]$.*

$$\text{Then } p_1 \leq \sqrt{3p_0p_2 + \delta^2} \leq \sqrt{3p_2} + \delta.$$

Proof. Let $\epsilon_i := \Pr[e_i = 1]$, and without loss of generality let

$$\delta = \epsilon_1 \geq \epsilon_2 \geq \dots \geq \epsilon_n$$

be given in descending order. Then

$$p_0 = \prod_{i=1}^n (1 - \epsilon_i), \quad p_1 = p_0 \cdot \left(\sum_{i=1}^n \frac{\epsilon_i}{1 - \epsilon_i} \right), \quad p_2 = p_0 \cdot \left(\sum_{i>j} \frac{\epsilon_i \epsilon_j}{(1 - \epsilon_i)(1 - \epsilon_j)} \right)$$

so that

$$\begin{aligned}
p_1^2 - 3p_0p_2 &= p_0^2 \cdot \left(\sum_{i,j=1}^n \frac{\epsilon_i \epsilon_j}{(1-\epsilon_i)(1-\epsilon_j)} - 3 \sum_{i>j} \frac{\epsilon_i \epsilon_j}{(1-\epsilon_i)(1-\epsilon_j)} \right) \\
&= p_0^2 \cdot \left(\sum_{i=1}^n \frac{\epsilon_i^2}{(1-\epsilon_i)^2} - \sum_{i>j} \frac{\epsilon_i \epsilon_j}{(1-\epsilon_i)(1-\epsilon_j)} \right) \\
&\geq p_0^2 \cdot \left(\sum_{i=1}^n \frac{\epsilon_i^2}{(1-\epsilon_i)^2} - \sum_{i=j+1} \frac{\epsilon_i \epsilon_j}{(1-\epsilon_i)(1-\epsilon_j)} \right) \\
&\geq p_0^2 \cdot \left(\sum_{i=1}^n \frac{\epsilon_i^2}{(1-\epsilon_i)^2} - \sum_{i=2}^n \frac{\epsilon_i^2}{(1-\epsilon_i)^2} \right) \\
&= p_0^2 \cdot \frac{\epsilon_1^2}{(1-\epsilon_1)^2} \leq \delta^2.
\end{aligned}$$

Hence, $p_1^2 - 3p_0p_2 \leq \delta^2$ and $p_1 \leq \sqrt{3p_0p_2 + \delta^2}$ as claimed. \square

We are now ready to prove Lemma 6.

Lemma 6. *Let $\mathbf{P} = (\text{Keygen}, \text{Encr}, \text{Decr})$ be a δ -correct, ϵ -injective rPKE with messages in \mathcal{M} and randomness in \mathcal{R} . Let $G : \mathcal{M} \rightarrow \mathcal{R}$ be a random oracle, so that $T(\mathbf{P}, G) := (\text{Keygen}, \text{Encr}_1, \text{Decr})$ is a derandomized version of \mathbf{P} . Let \mathcal{A} be a FFC adversary against $T(\mathbf{P}, G)$ which makes at most q queries at depth d to G and returns a list of at most q_{dec} ciphertexts. Then*

$$\text{Adv}_{T(\mathbf{P}, G)}^{\text{FFC}}(\mathcal{A}) \leq ((4d+1)\delta + \sqrt{3\epsilon}) \cdot (q + q_{\text{dec}}) + \epsilon.$$

Proof. Essentially, the idea is as follows: the adversary gets an advantage of about $4dq\delta$ from querying G in search of failing ciphertexts, and at most $q_{\text{dec}}(\delta + 3\sqrt{\epsilon})$ from guessing blindly. The latter term comes from considering ways that some blind guess could be a failing ciphertext: if it is the encryption of one message, then δ is large, and if it is possibly the encryption of more than one message (e.g., as a general “encryption failed” output), then ϵ is large. We will formalize this in what follows.

Generate a keypair $(\text{pk}, \text{sk}) \leftarrow \text{Keygen}()$ and oracle $G \leftarrow^{\$} \mathcal{R}^{\mathcal{M}}$. Let

$$Y_m := \{r : \text{Decr}(\text{sk}, \text{Encr}(\text{pk}, m, r)) = m\}$$

be the set of coins such that decryption of m will succeed. Let $G'(m) := G(m)$ if $G(m) \in Y_m$, $G'(m) \leftarrow^{\$} \mathcal{R}$ if $Y_m = \emptyset$, and $G'(m) \leftarrow^{\$} Y_m$ otherwise. Thus G' is uniformly random in the space \mathcal{G} of oracles where decryption succeeds if possible. Moreover, G' is independent of the behavior of messages and ciphertexts for $T(\mathbf{P}, G)$ which do not decrypt correctly.

Now, fix (sk, pk) and G' and let

$$\delta' := \max_{m \in \mathcal{M}} \Pr[\text{Decr}(\text{sk}, \text{Encr}(\text{pk}, m)) \neq m]$$

be the failure probability for this keypair. Let DbfFail be the event that some ciphertext c is the encryption of two messages m_1 and m_2 such that $\text{Decr}(\text{sk}, c) \notin \{m_1, m_2\}$. We define $\epsilon' := \Pr[\text{DbfFail}]$. Both δ' and ϵ' are independent of G' . In addition, let Fail be the event that \mathcal{A} wins the FFC game (see Definition 3), and $\text{Ev} := \text{Fail} \wedge \neg \text{DbfFail}$. By Lemma 1, it holds that

$$\left| \sqrt{\Pr[\text{Ev} : \mathcal{A}^G(\text{pk})]} - \sqrt{\Pr[\text{Ev} : \mathcal{A}^{G'}(\text{pk})]} \right| \leq 2d\sqrt{P_{\text{guess}}}.$$

Since, conditioned on G' , $G(m) \neq G'(m)$ at each m with probability at most δ' and there are q/d guesses (in expectation), it holds furthermore that

$$2d\sqrt{P_{\text{guess}}} \leq \sqrt{4d^2 P_{\text{guess}}} \leq \sqrt{4dq\delta'}.$$

Next we define for a ciphertext c ,

$$p_1(c) := \Pr[\exists \text{ unique } m \in \mathcal{M} : c = \text{Encr}(\text{pk}, m, G(m)) \wedge \text{Decr}(\text{sk}, c) \neq m].$$

(It is important to note that if m exists but is not unique, then DbfFail occurs.) Furthermore, let $p_1 := \max_c(p_1(c))$. Since $p_1(c)$ and p_1 are independent of G' , we have

$$\Pr[\text{Ev} : \mathcal{A}^{G'}(\text{pk})] \leq q_{\text{dec}} \cdot p_1.$$

By Lemma 7, $p_1 \leq \delta' + \sqrt{3\epsilon'}$. Plugging this in and applying the Cauchy-Schwarz corollary $\sqrt{ab} + \sqrt{cd} \leq \sqrt{(a+c)(b+d)}$ gives

$$\begin{aligned} \sqrt{\Pr[\text{Ev} : \mathcal{A}^G(\text{pk})]} &\leq \sqrt{4dq\delta'} + \sqrt{q_{\text{dec}} \cdot (\delta' + \sqrt{3\epsilon'})} \\ &\leq \sqrt{((4d+1)\delta' + \sqrt{3\epsilon'}) \cdot (q + q_{\text{dec}})}. \end{aligned}$$

Finally, by definition of correctness and injectivity of an rPKE (see Definition 5 and 6, respectively), it holds that $\delta = \mathbb{E}[\delta']$ and $\epsilon \leq \mathbb{E}[\epsilon']$. By Jensen's inequality, it holds furthermore that $\sqrt{\epsilon} \leq \mathbb{E}[\sqrt{\epsilon'}]$. Hence,

$$\begin{aligned} \text{Adv}_{T(\mathcal{P}, G)}^{\text{FFC}}(\mathcal{A}) &\leq \mathbb{E} \left[\Pr[\text{Ev} : \mathcal{A}^G(\text{pk})] : (\text{pk}, \text{sk}) \leftarrow \text{Keygen}(); G' \stackrel{\$}{\leftarrow} \mathcal{G} \right] + \epsilon \\ &\leq ((4d+1)\delta + \sqrt{3\epsilon}) \cdot (q + q_{\text{dec}}) + \epsilon \end{aligned}$$

as claimed. \square

F Proof of Theorem 4

Theorem 4 (Implicit \rightarrow explicit with key confirmation). *Let \mathcal{P} be an ϵ -injective dPKE. Consider the KEM $\mathcal{K}_1 := U_m^\perp(C(\mathcal{P}, H_t, \tau), H_s)$ obtained from \mathcal{P} applying the C-transform with random oracle $H_t : \mathcal{M} \rightarrow \{0, 1\}^\tau$ and the U_m^\perp -transform with independent random oracle $H_s : \mathcal{M} \rightarrow \{0, 1\}^s$. Let $\mathcal{K}_2 :=$*

$U_m^\chi(\mathsf{P}, \mathsf{F}, H)$ be the KEM obtained from P applying the U_m^χ -transform with random oracle $H : \mathcal{M} \rightarrow \{0, 1\}^{\varsigma+\tau}$.

If \mathcal{A} is an IND-CCA-adversary against K_1 which makes q_{dec} decapsulation queries, then it is also an IND-CCA-adversary against P_2 and there is a PRF-adversary \mathcal{B} against F which uses about the same time and resources as \mathcal{A} , such that:

$$\text{Adv}_{\mathsf{K}_1}^{\text{IND-CCA}}(\mathcal{A}) \leq 2 \cdot \text{Adv}_{\mathsf{K}_2}^{\text{IND-CCA}}(\mathcal{A}) + \frac{q_{\text{dec}}}{2^{\tau-1}} + 2 \cdot \text{Adv}_{\mathsf{F}}^{\text{PRF}}(\mathcal{B}) + 2\epsilon.$$

Proof. The proof is by a series of games. Let w_i be the probability that \mathcal{A} wins Game i .

Game 0 (IND-CCA). *This is the IND-CCA game against K_1 .*

Game 1 (Modify decapsulation with random function). *In Game 1, the simulator instantiates a fresh random function R , and modifies the decapsulation oracle D to the oracle D' shown in Figure 8.*

$D((c, t)):$	$D'((c, t)):$
1 if $(c, t) = (c^*, t^*)$: return \perp	1 if $c = c^*$: return \perp
2 $m' \leftarrow \text{Decr}(\text{sk}, c)$	2 $m' \leftarrow \text{Decr}(\text{sk}, c)$
3 if $m' = \perp$:	3 if $m' = \perp$: $(k, t') \leftarrow R(c)$
4 return \perp	4 else if $\text{Encr}(\text{pk}, m') \neq c$: $(k, t') \leftarrow R(c)$
5 else if $(\text{Encr}(\text{pk}, m'), H_t(m')) \neq (c, t)$:	5 else: $(k, t') \leftarrow (H_s(m'), H_t(m'))$
6 return \perp	6 if $t' \neq t$: return \perp
7 else: return $H_s(m')$	7 else: return k

Fig. 8. Decapsulation oracles for Game 0 and Game 1.

We analyze the difference between D and D' as follows.

- If $c = c^*$, then D' returns \perp . If $\text{Encr}(\text{pk}, \cdot)$ is injective, then so does D .
- Otherwise, if $\text{Encr}(\text{pk}, m') = c$, then both D and D' return $H_s(m')$ if $H_t(m') = t$, and \perp otherwise.
- Otherwise, D returns \perp , and so does D' unless t matches t' . Since R is random and is only used for this purpose, this happens with probability at most $q_{\text{dec}}/2^\tau$.

The difference in \mathcal{A} 's view is bound by the probability that D' acts different than D . So overall $|w_1 - w_0| \leq q_{\text{dec}}/2^\tau + \epsilon$.

Game 2 (Use PRF instead of R). *In Game 2, the simulator replaces $R(c)$ by $\mathsf{F}(\text{prfk}, \cdot)$ with a random prf key prfk .*

The difference in probability of any adversary \mathcal{A} between winning Game 1 and Game 2 is exactly the PRF-advantage of an adversary \mathcal{B} that works exactly as in the analysis of Game 1 in the proof of Theorem 2. Hence it holds that

$$|w_2 - w_1| = \text{Adv}_{\mathbb{F}}^{\text{PRF}}(\mathcal{B})$$

Game 3 (Redirect to $U_m^{\mathcal{X}}(\mathbb{P}, \mathbb{F}, H)$). *Game 3 is refactored so that it simulates the IND-CCA experiment for $\mathbb{K}_2 = U_m^{\mathcal{X}}(\mathbb{P}, \mathbb{F}, H)$ for the case $b = 0$ (correct key), as follows:*

- *Hash redirection.* The oracles H_s resp. H_t used for C resp. U^{\perp} are redirected to the first ς resp. last t bits of the hash function H of $U_m^{\mathcal{X}}(\mathbb{P}, \mathbb{F}, H)$.
- *The simulator creates a challenge ciphertext c with shared secret (k_s, k_t) . This is given to \mathcal{A} as a challenge ciphertext (c, k_t) with shared secret $k = k_s$.*
- *The decapsulation oracle D' from Game 2 is changed to use the $U_m^{\mathcal{X}}$ decapsulation oracle internally, as shown in Figure 9. It is called D'' .*

$D''((c, t)):$

- 1 $r \leftarrow \text{Decaps}_{\mathbb{K}_2}(c)$
 - 2 if $r = \perp$: return \perp
 - 3 parse r as (k, t')
 - 4 if $t' \neq t$: return \perp
 - 5 else: return k
-

Fig. 9. Decapsulation oracle for Game 3.

All these steps do not change \mathcal{A} 's view compared to Game 2, so $w_3 = w_2$.

Game 4 (Redirect to $U_m^{\mathcal{X}}(\mathbb{P}, \mathbb{F}, H)$ with random keys). *Game 4 is the same as Game 3 except that it now simulates the $b = 1$ (random key) case of the IND-CCA experiment against \mathbb{K}_2 , i.e., it always sets $k \xleftarrow{\$} \{0, 1\}^{\varsigma}$.*

Distinguishing Game 3 from Game 4 is exactly the IND-CCA experiment for \mathbb{K}_2 . Hence,

$$|w_4 - w_3| = \text{Adv}_{\mathbb{K}_2}^{\text{IND-CCA}}(\mathcal{A})$$

Note that \mathcal{A} 's goal is still to learn the bit b of the IND-CCA game against \mathbb{K}_1 . But now the shared secret k is always random, and thereby independent of b . Hence, the adversary has no information about b and so $w_4 = \frac{1}{2}$.

Summing up the differences in winning probability from all the games we get

$$\left| w_0 - \frac{1}{2} \right| \leq \text{Adv}_{\mathbb{K}_2}^{\text{IND-CCA}}(\mathcal{A}) + \frac{q_{\text{dec}}}{2^{\tau}} + \text{Adv}_{\mathbb{F}}^{\text{PRF}}(\mathcal{B})$$

and $\text{Adv}_{\mathbb{K}_1}^{\text{IND-CCA}}(\mathcal{A})$ is at most twice this value. This completes the proof. \square