Post-Quantum Security of Key Encapsulation Mechanism against CCA Attacks with a Single Decapsulation Query

Haodong Jiang[1] * Zhi Ma * Zhenfeng Zhang †

*Henan Key Laboratory of Network Cryptography Technology

†Institute of Software, Chinese Academy of Sciences

Dec., 2023

[1] Presented by **Yiting Liu**

Overview

Background

Diffie-Hellman Key Exchange A fundamental and elegant cryptographic scheme.

Current Application Ephemeral key establishment in TLS, Signal, etc..

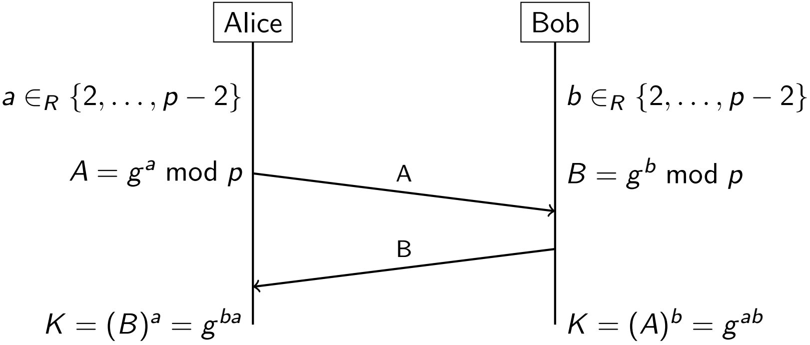

Background

Diffie-Hellman (DH) key exchange A fundamental and elegant cryptographic scheme.

Current Application Ephemeral key establishment in TLS, Signal, etc..

Shor's algorithm

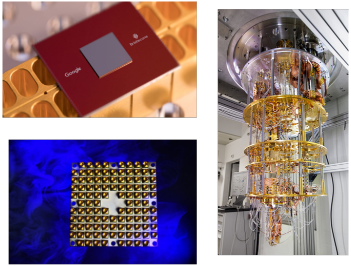

Rapid advance in quantum computing

PQC and NIST's Standardization

Post-Quantum Cryptography (PQC) *classical* cryptosystems that remain secure in the presence of a quantum adversary

NIST's PQC Standardization PKE, Digital signatures and **KEM**

$$\mathrm{KEM} = (Gen, Encap, Decaps)$$

Alice | Bob

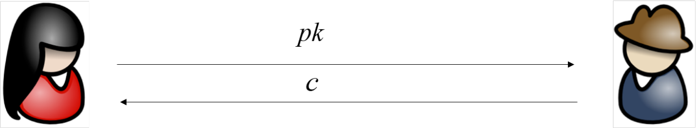

$$Gen(1^{\lambda}) \rightarrow (pk, sk)$$

$$Decaps(sk, c) \rightarrow K$$

$$Encap(pk) \rightarrow (K, c)$$

PQC and NIST's Standardization

Post-Quantum Cryptography (PQC) *classical* cryptosystems that remain secure in the presence of a quantum adversary

NIST's PQC Standardization PKE, Digital signatures and **KEM**

- August 24, 2023, NIST posted the first KEM standard draft (Kyber) FIPS-203.
- Kyber is a lattice-based KEM with IND-CCA security.

IND-CCA security

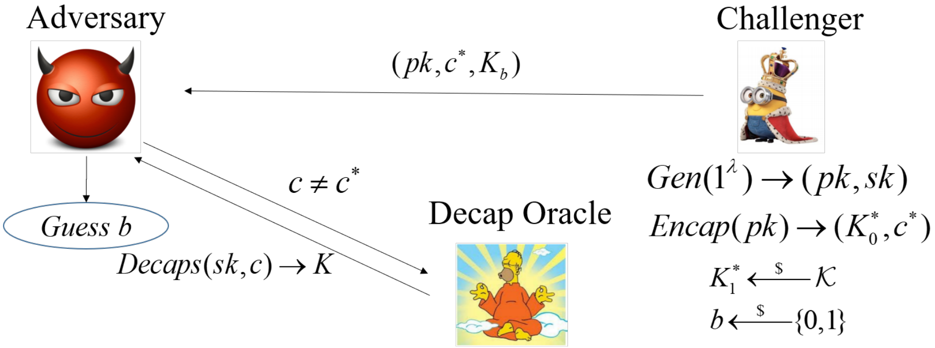

$$\mathsf{Adv}_{\mathrm{KEM}}^{\mathrm{IND\text{-}CCA}}(\mathcal{A}) := \left|\Pr[\mathrm{IND\text{-}CCA}_{\mathrm{KEM}}^{\mathcal{A}} = 1] - 1/2\right|$$

Generic constructions of an IND-CCA-secure KEM

FO-like generic constructions: weakly-secure PKE $\Rightarrow$ CCA-secure KEM

Gen'

$(pk, sk) \leftarrow Gen$

$s \xleftarrow{\$} \mathcal{M}$

$sk' := (sk, s)$

return (pk, sk')

$Encaps(pk)$

$m \xleftarrow{\$} \mathcal{M}$

$c = Enc(pk, m; G(m))$

$K := H(m, c)$

return (K, c)

$Decaps(sk', c)$

Parse $sk' = (sk, s)$

$m' := Dec(sk, c)$

if $Enc(pk, m'; G(m')) = c$

  return $K := H(m', c)$

else return

  $K := H(s, c)$

Figure: IND-CCA-secure KEM$=\mathrm{FO}^{\not\perp}$[PKE,G,H]

Generic constructions of an IND-CCA-secure KEM

FO-like generic constructions: weakly-secure PKE $\Rightarrow$ CCA-secure KEM

- Re-encryption in decapsulation makes it an expensive operation. As shown by [HV22], when re-encryption is removed, there will be a 2.17X and 6.11X speedup over decapsulation in Kyber and FrodoKEM respectively.
- The re-encryption makes the KEM more vulnerable to side-channel attacks and almost all the NIST-PQC Round-3 KEMs are affected [MeI22].
- The side-channel protection of re-encryption will significantly increase deployment costs and thus complicate the integration of NIST-PQC KEMs [MeI22].

Diffie-Hellman ⇒ KEM

- For ephemeral key establishment, one has to move the current DH key-exchange to post-quantum KEMs.
- IND-1CCA security is required for such a substitutive KEM in post-quantum TLS 1.3 [HV22], KEM-TLS [SSW20], post-quantum Signal [BFGJS22] and post-quantum Noise [ADHSW22].
- IND-1CCA security is the same as the IND-CCA security except that the adversary is restrictive to make at most one *single* decapsulation query.
- Obviously, IND-1CCA security is implied by the IND-CCA security. However, the current IND-CCA-secure KEMs require re-encryption.

Designing a dedicated IND-1CCA-secure KEM without re-encryption was taken as an open problem raised by Schwabe, Stebila and Wiggers [SSW20].

Huguenin-Dumittan and Vaudenay's work [HV22]

Huguenin-Dumittan and Vaudenay shows transforms T_{CH} and T_H can turn a CPA-secure PKE into an IND-1CCA-secure KEM.

Gen	$Encaps(pk)$	$Decaps(sk, (c, tag))$
$(pk, sk) \leftarrow Gen'$	$m \leftarrow_{\$} \mathcal{M}$	$m' := Dec'(sk, c)$
return (pk, sk)	$c \leftarrow Enc'(pk, m)$	**if** $H'(m', c) = tag // T_{CH}$
	$tag = H'(m, c) // T_{CH}$	**if** $m' = \perp // T_H$
	$K := H(m) // T_{CH}$	**return** $\perp$
	return $(K, (c, tag)) // T_{CH}$	**else return** $K := H(m')$
	$K := H(m, c) // T_H$	
	return $(K, c) // T_H$	

Figure: $\text{KEM}_{CH} = T_{CH}[\mathrm{PKE}', H]$ and $\text{KEM}_H = T_H[\mathrm{PKE}', H]$

Quantum random oracle model

- The constructions KEM_{CH} and KEM_{H} are based on an idealized model called random oracle model (ROM), where a hash function is idealized to be a publicly accessible random oracle (RO).
- In post-quantum setting, quantum adversary can execute hash functions (the instantiation of **RO**) on an arbitrary superposition of inputs.
- Therefore, Boneh et al. [BDF+11] argued that to prove post-quantum security one needs to prove security in the quantum random oracle model (QROM), where the adversary can query the RO with quantum state.

Huguenin-Dumittan and Vaudenay's work [HV22]

- The security of T_{CH} was proved in the ROM with tightness $\epsilon_R \approx O(1/q)\epsilon_{\mathcal{A}}$, and in the QROM with tightness $\epsilon_R \approx O(1/q^3)\epsilon_{\mathcal{A}}^2$.
- The security of T_H was proved in the ROM with tightness $\epsilon_R \approx O(1/q^3)\epsilon_{\mathcal{A}}$. The QROM proof of T_H was left open.
- Both T_{CH} and T_H do not require re-encryption. But, compared with T_{CH}, T_H does not need the key confirmation and thus will not lead to ciphertext expansion.

Our results

- First, we prove the security of T_H and its implicit variant T_{RH} in both ROM and QROM. T_{RH} is the same as the T_H except that in decapsulation a pseudorandom value $H(\star, c)$ is returned instead of an explicit $\perp$ for an invalid ciphertext c such that $Dec(sk, c) = \perp$.

Gen	$Encaps(pk)$	$Decaps(sk, c)$
$(pk, sk) \leftarrow Gen'$	$m \leftarrow_\$ \mathcal{M}$	$m' := Dec'(sk, c)$
return (pk, sk)	$c \leftarrow Enc'(pk, m)$	**if** $m' = \perp$
	$K := H(m, c)$	**return** $\perp$ $\quad // T_H$
	return (K, c)	**return** $K := H(\star, c)$ $\quad // T_{RH}$
		else return $K := H(m', c)$

Figure: $\text{KEM}_H = T_H[\mathrm{PKE}', H]$ and $\text{KEM}_{RH} = T_{RH}[\mathrm{PKE}', H]$

Remarks on T_{RH}

- T_{RH} is essentially the construction $U^{\not\perp}$ in [HHK17], except that the secret seed s in decapsulation is replaced by a public value $\star$ ($\star$ can be any fixed message).
- In fact, our proof can work for both secret seed and public value thanks to the newly introduced decapsulation simulation technique, while the current IND-CCA proofs for implicit FO-KEMs [HHK17, JZC+18] can only work for secret seed.
- We choose to replace secret seed by public value since it reduces the secret key size and makes the construction more concise.
- Moreover, from a high-assurance implementation (i.e., side-channel protected) point of view, public value is also preferable to secure seed [Sch22].

The tightness of the reduction

Table: Reduction tightness in the ROM/QROM.

Transformation	Reduction tightness	Ciphertext expansion	Re-encryption	ROM or QROM
FO [HHK17]	$\epsilon_R \approx \epsilon_{\mathcal{A}}$	N	Y	ROM
T_{CH} [HV22]	$\epsilon_R \approx O(1/q)\epsilon_{\mathcal{A}}$	Y	N	ROM
T_H [HV22]	$\epsilon_R \approx O(1/q^3)\epsilon_{\mathcal{A}}$	N	N	ROM
Our T_{RH} and T_H	$\epsilon_R \approx O(1/q)\epsilon_{\mathcal{A}}$	N	N	ROM
FO [JZM19,BHH+19]	$\epsilon_R \approx O(1/q)\epsilon_{\mathcal{A}}^2$	N	Y	QROM
T_{CH} [HV22]	$\epsilon_R \approx O(1/q^3)\epsilon_{\mathcal{A}}^2$	Y	N	QROM
Our T_{RH} and T_H	$\epsilon_R \approx O(1/q^2)\epsilon_{\mathcal{A}}^2$	N	N	QROM

The tightness of the reduction

- Then, for T_H, T_{RH} and T_{CH}, we show that if the underlying PKE meets malleability property, a $O(1/q)$ ($O(1/q^2)$, resp.) loss is unavoidable in the ROM (QROM, resp.).
- That is, our ROM reduction is optimal in general. Roughly speaking, the malleability property says that an adversary can efficiently transform a ciphertext into another ciphertext which decrypts to a related plaintext.
- In particular, such a malleability property is met by real-world PKE schemes, e.g., ElGamal, FrodoKEM.PKE, Kyber.PKE, etc.

Relations among notions of CCA security for KEM

- Finally, we compare the relative strengths of IND-1-CCA and IND-CCA in the ROM and QROM. For each pair of notions A, B $\in$\{IND-1-CCA ROM, IND-CCA ROM, IND-1-CCA QROM, IND-CCA QROM\}, we show either an implication or a separation, so that no relation remains open.

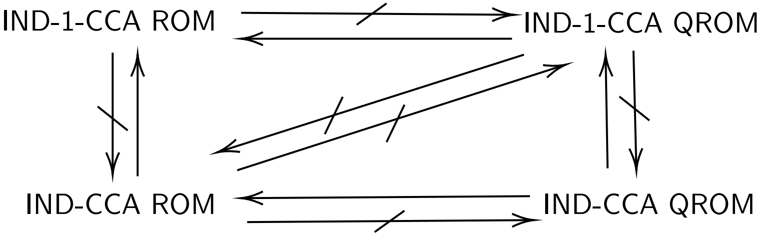

Figure: The relations among notions of security for KEM.

Construction and reduction

Gen	$Encaps(pk)$	$Decaps(sk, c)$
$(pk, sk) \leftarrow Gen'$	$m \leftarrow_{\$} \mathcal{M}$	$m' := Dec'(sk, c)$
return (pk, sk)	$c \leftarrow Enc'(pk, m)$	**if** $m' = \bot$
	$K := H(m, c)$	**return** $\bot$ $//T_H$
	return (K, c)	**return** $K := H(\star, c)$ $//T_{RH}$
		else return $K := H(m', c)$

Figure: $\text{KEM}_H = T_H[\mathrm{PKE}', H]$ and $\text{KEM}_{RH} = T_{RH}[\mathrm{PKE}', H]$

Main theorem

Theorem 3.1 (QROM security of T_{RH}).

For any adversary $\mathcal{B}$ against the IND-1-CCA security of KEMs, issuing at most one single (classical) query to the decapsulation oracle and at most q_H queries to the quantum random oracle H, there exists an IND-CPA adversary $\mathcal{D}$ against PKE′ such that

$$\mathbf{Adv}^{\text{IND-1-CCA}}_{\text{KEM}_{RH}}(\mathcal{B}) \leq 6(q_H+1)\sqrt{4\mathbf{Adv}^{\text{IND}-\text{CPA}}_{\text{PKE}'}(\mathcal{D}) + 2(q_H+1)^2/|\mathcal{M}| + 1/|\mathcal{K}|}$$

$$\mathbf{Adv}^{\text{IND-1-CCA}}_{\text{KEM}_{H}}(\mathcal{B}') \leq \mathbf{Adv}^{\text{IND-1-CCA}}_{\text{KEM}_{RH}}(\mathcal{B}) + \epsilon_{\text{coll}},$$

where $\text{Time}(\mathcal{D}) \approx \text{Time}(\mathcal{B}) + O(q_H^2)$, $\text{Time}(\mathcal{B}') \approx \text{Time}(\mathcal{B})$, ϵ_{coll} *is an advantage bound of an algorithm searching a collision of the random oracle H with q_H queries. In particular,* $\epsilon_{\text{coll}} = q_H^2/|\mathcal{K}|$ *in the ROM, and* $\epsilon_{\text{coll}} = q_H^3/|\mathcal{K}|$ *in the QROM.*

Proof Skeleton of the Main Theorem

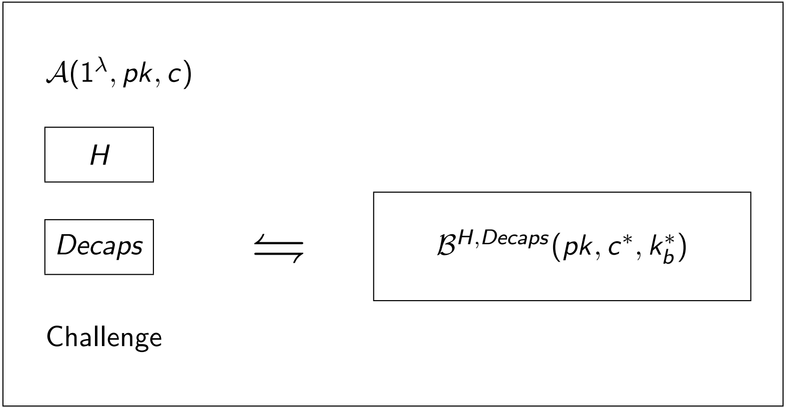

The simulation of the decapsulation oracle

- Re-encryption is the core feature of FO-like CCA-KEMs, which guarantees that only specific valid ciphertexts can be correctly decapsulated, and thus makes the decapsulation simulation in the ROM/QROM proof easy.
- Removing re-encryption makes the current decapsulation simulation for FO-like CCA-KEMs incompatible with the KEMs in this paper.
- For a valid ciphertext $\bar{c}$ such that $(Dec(sk, \bar{c}) = \bar{m} \neq \perp)$, the decapsulation returns $H(\bar{m}, \bar{c})$.
- If we reprogram $H(\bar{m}, \bar{c})$ to a random $\bar{k}$, we can simulate the decapsulation of $\bar{c}$ using $\bar{k}$ without knowledge of sk.
- To guarantee the consistency between the outputs of H and the simulated decapsulation, one needs to correctly guess when the adversary makes a query $(\bar{m}, \bar{c})$ to H, and perform a reprogram at that time. In the ROM, a randomly guess is correct with probability $1/q$.

The simulation of the decapsulation oracle

- In the QROM, due to adversary's superposition RO-query, it is hard to define when the adversary makes a query $(\bar{m}, \bar{c})$. We find that the consistency between H and the simulated decapsulation can be guaranteed if the predicate $Decap(sk, \bar{c}) = H(\bar{m}, \bar{c})$ is satisfied.
- Don, Fehr, Majenz, and Schaffner [DFMS19, DFM20] showed that a random measure-and-reprogram can keep the predicate satisfied with a high probability.
- However, the measure-and-reprogram in [DFMS19, DFM20] cannot be directly applied to our case. This is due to the fact that the random measure in [DFMS19, DFM20] is performed for all the H-queries while in our case there is an implicit (classical) H-query used in the real decapsulation that will be removed in the simulated decapsulation and thus can not be measured.
- Extending the measure-and-reprogram technique in [DFMS19, DFM20], we derive a variant of measure-and-reprogram, which is suitable for our case. With this new measure-and-reprogram, the QROM adversary can accept the simulations with probability at least $O(1/q^2)$.

Measure-and-reprogram and our extension

Standard Measure-and-reprogram [DFM20]: $\Pr_H[V(x, H(x), z) = 1 : (x, z) \leftarrow A^{|H\rangle}] \leq$

$$O(q^2) \Pr_{H,\Theta}[V(x, \Theta, z) = 1 : (x, z) \leftarrow S^A]$$

Our (Single-Classical-Query) version: $\Pr_H[V(x, H(x), z) = 1 : (x, z) \leftarrow A^{|H\rangle}] \leq$

$$O(q^2) \Pr_{H,\Theta}[V(x, \Theta, z) = 1 : (x, z) \leftarrow S_1^A] + O(q^2) \Pr_{H,\Theta}[V(x, \Theta, z) = 1 : (x, z) \leftarrow S_2^A]$$

- A^H an arbitrary q-query quantum algorithm
- S^A an algorithm that randomly measure and reprogram on all A's H-queries
- S_1^A an algorithm that randomly measure and reprogram on all A's H-queries except for one specific classical H-query x
- S_2^A an algorithm that randomly measure and reprogram on A's H-queries after A makes the specific classical H-query x

The simulation of the decapsulation oracle

Lemma 3.2 ((Single-Classical-Query) informal Measure-and-reprogram).

Let $A^{|H\rangle}$ be an arbitrary oracle quantum algorithm that makes q queries to H, and outputs some classical x and a (possibly quantum) output z. In particular, A's i^*-th query input state is exactly x. Let $S_1^A(\Theta)$ be an oracle algorithm that answers A's i^*-th query with Θ, randomly measures and reprograms on A's other queries. Finally, $S_1^A(\Theta)$ returns A's output. Let $S_2^A(\Theta)$ be an oracle algorithm that only randomly measures and reprograms on A's j^*-th queries ($j \geq i^*$). Finally, $S_2^A(\Theta)$ returns A's output. Thus, for any $x_0 \in X$, $i^* \in \{1, \cdots, q\}$ and any predicate V:

$$\Pr_H[x = x_0 \wedge V(x, H(x), z) = 1 : (x, z) \leftarrow A^{|H\rangle}] \leq 2(2q-1)^2 \Pr_{H,\Theta}[x = x_0 \wedge V(x, \Theta, z) = 1 : (x, z) \leftarrow S_1^A] + 8q^2 \Pr_{H,\Theta}[x = x_0 \wedge V(x, \Theta, z) = 1 : (x, z) \leftarrow S_2^A],$$

The embedding of the hard instance

- We use the oneway-to-hiding (O2H) technique to embed the instance of the underlying IND-CPA-security experiment.

Lemma 3.3 ((Adapted) Double-sided O2H [BHH+19]).

Let G, H : be oracles such that $\forall x \neq x^*$. $G(x) = H(x)$. Let z be a random bitstring. Let A be quantum oracle algorithm that makes at most q queries (not necessarily unitary). Then, there is an another double-sided oracle algorithm $B^{|G\rangle,|H\rangle}(z)$ such that B runs in about the same amount of time as A, and

$$\left|\Pr[1 \leftarrow A^{|H\rangle}(z)] - \Pr[1 \leftarrow A^{|G\rangle}(z)]\right| \leq 2\sqrt{\Pr[x^* = x' : x' \leftarrow B^{|G\rangle,|H\rangle}(z)]}.$$

The embedding of the hard instance

Lemma 3.4 (Search in Double-sided Oracle).

Let G, H : be oracles such that $\forall x \neq x^* \; G(x) = H(x)$. Let z be a random bitstring. Let A be quantum oracle algorithm that makes at most q queries (not necessarily unitary). Let $B^{|G\rangle,|H\rangle}(z)$ be a double-sided oracle algorithm defined in Lemma 3.3. Let $C^{|H\rangle}(z)$ be an oracle algorithm that picks $i \leftarrow_{\$} \{1, 2, \ldots, q\}$, runs $A^{|H\rangle}(z)$ until (just before) the i-th query, measures the query input registers in the computational basis, and outputs the measurement outcome. Thus, we have

$$\Pr[x^* = x' : x' \leftarrow B^{|G\rangle,|H\rangle}(z)] \leq q^2 \Pr[x^* = x' : x' \leftarrow C^{|H\rangle}(z)].$$

In particular, if $\mathcal{X} = \mathcal{X}_1 \times \mathcal{X}_2$, $x^* = (x_1^*, x_2^*)$, x_1^* is uniform and independent of H and z, then we further have $\Pr[x^* = x' : x' \leftarrow B^{|G\rangle,|H\rangle}(z)] \leq q^2 / |\mathcal{X}_1|$.

Conclusion

1. An IND-1-CCA KEM is sufficient to replace Diffie-Hellman in the post-quantum migration of the widely-deployed protocols, such as TLS 1.3, Signal and Noise.
2. Our results show that IND-1-CCA-secure KEMs can be constructed in the ROM and QROM without re-encryption and cipher-expansion.
3. Compared with IND-CCA-secure KEMs based on FO transform, the IND-1-CCA-secure KEMs based on T_H and T_{RH} do not require the re-encryption in decapsulation. The re-encryption is highly vulnerable to attacks and its side-channel protection will significantly increase deployment costs.
4. Thus, from a practical point of view, removing the re-encryption of FO-like KEMs will improve the performance of embedded side-channel secure implementations.
5. Therefore, according to our results, one can easily transform Kyber.PKE into an IND-1-CCA-secure KEM without re-encryption and cipher-expansion, and then establish post-quantum-secure variants of TLS 1.3, Signal and Noise with better performance in the embedded implementation.

Thanks for your attention!

hdjiang13@gmail.com

References

HV22 Huguenin-Dumittan, L., Vaudenay, S.: On IND-qCCA security in the ROM and its applications - CPA security is sufficient for TLS 1.3.

Mel22 Melissa Azouaoui et al. Surviving the fo-calypse: Securing pqc implementations in practice. RWC 2022 (2022)

SSW20 Schwabe, P., Stebila, D., Wiggers, T.: Post-quantum TLS without handshake signatures.

BFGJS22 Brendel, J., Fiedler, R., Günther, F., Janson, C., Stebila, D.: Post-quantum asynchronous deniable key exchange and the signal handshake

Sch22 Schneider, T.: Implicit rejection in kyber. NIST pqc-forum (2022)

BDF+11 Boneh, D., Dagdelen, O., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random oracles in a quantum world

References

HHK17 Dennis Hofheinz, Kathrin Hövelmanns and Eike Kiltz, A modular analysis of the Fujisaki-Okamoto transformation

JZC+18 Haodong Jiang et al., IND-CCA-secure Key Encapsulation Mechanism in the Quantum Random Oracle Model, Revisited

DFM19 Don, J., Fehr, S., Majenz, C., Schaffner, C.: Security of the Fiat-Shamir transformation in the quantum random-oracle model

DFM20 Don, J., Fehr, S., Majenz, C.: The measure-and-reprogram technique 2.0: Multi-round fiat-shamir and more

BHH+19 Bindel, N., Hamburg, M., Hövelmanns, K., Hülsing, A., Persichetti, E.: Tighter proofs of CCA security in the quantum random oracle model