On the non-tightness of measurement-based reductions for key encapsulation mechanism in the quantum random oracle model

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

[†]TCA Laboratory, Institute of Software, Chinese Academy of Sciences

[*]Chinese State Key Laboratory of Mathematical Engineering and Advanced Computing

hdjiang13@gmail.com

November 30, 2021

Overview

Background

Main Contribution

Techniques

Conclusion

Background

Public Key Cryptography Public Key Encryption (PKE), Digital Signatures, and Key Encapsulation Mechanism (KEM)

Current Deployment Diffie-Hellman key exchange, the RSA cryptosystem, and elliptic curve cryptosystems

Shor's algorithm

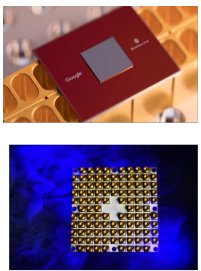

Rapid advance in quantum computing

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

- Dec. 2016 – Call for Proposals
- Dec. 2017 – Round-1-submissions (35/69)
- Jan. 2019 – Round-2-submissions (17/26)
- Aug. 2020 – Round-3-submissions (Finalists: 4/7, Alternates: 5/8)

KEM and IND-CCA security

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

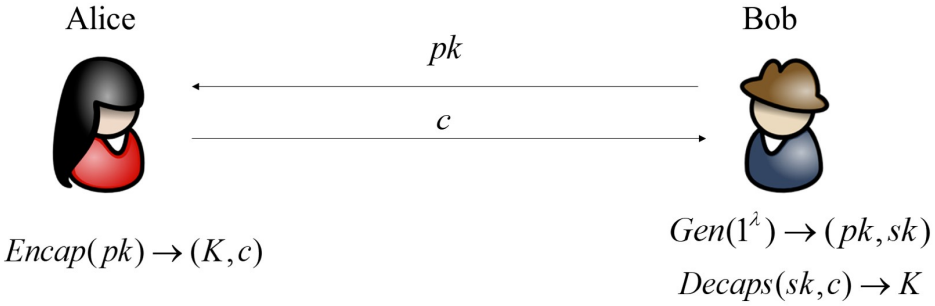

KEM and IND-CCA security

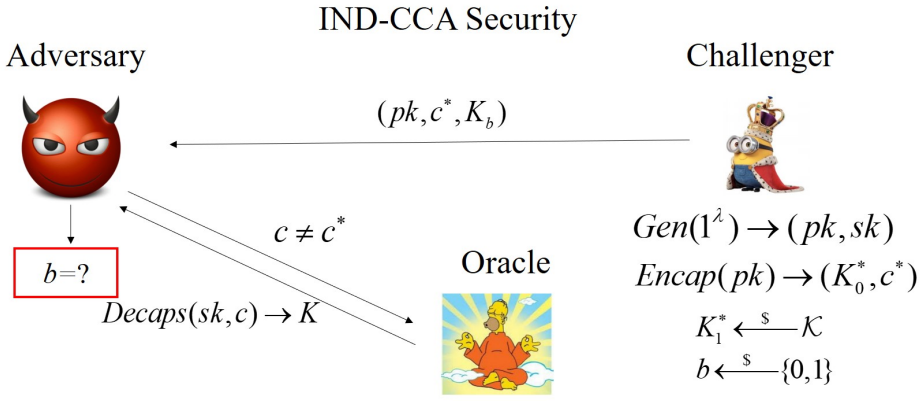

Advantage Function: $\mathsf{Adv}^{\text{IND-CCA}}_{\text{KEM}}(\mathcal{A}) := \left|\Pr[\text{IND-CCA}^{\mathcal{A}}_{\text{KEM}} = 1] - 1/2\right|$

Generic Construction

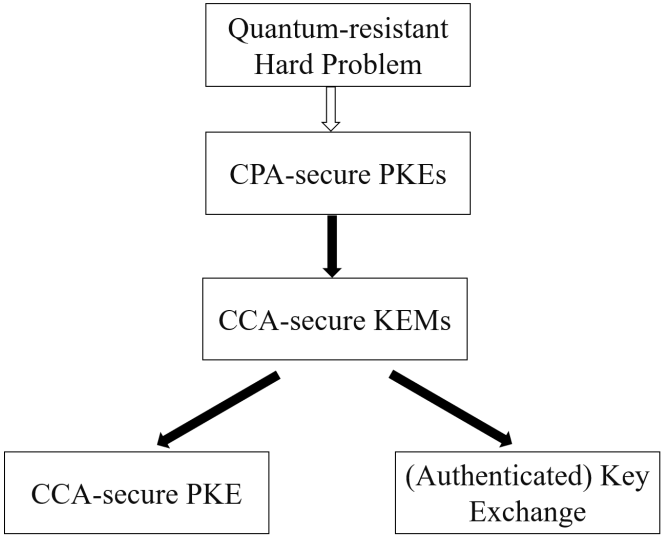

Generic Construction

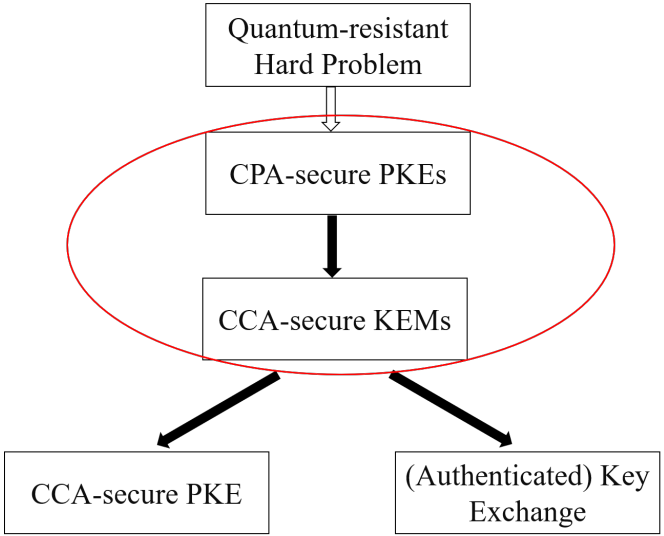

FO-like generic constructions

All currently known CPA $\Rightarrow$ CCA PKE/KEMs are Fujisaki-Okamoto(FO)-like and can be classified based on the underlying assumptions,

1. the variants of FO: $\mathrm{FO}^{\not\perp}$, $\mathrm{FO}^{\perp}$, $\mathrm{FO}^{\not\perp}_m$, $\mathrm{FO}^{\perp}_m$, $\mathrm{QFO}^{\not\perp}_m$ and $\mathrm{QFO}^{\perp}_m$ [1]
2. the variants of REACT/GEM: $\mathrm{U}^{\not\perp}$, $\mathrm{U}^{\perp}$, $\mathrm{U}^{\not\perp}_m$, $\mathrm{U}^{\perp}_m$, $\mathrm{QU}^{\not\perp}_m$ and $\mathrm{QU}^{\perp}_m$ [2]

Note: FO-like generic constructions are widely used in the NIST Round-3 KEM Candidates.

[1] m (without m) means $K = H(m)$ ($K = H(m, c)$), Q means an additional length-preserving hash [TU16] is added into the ciphertext, and $\not\perp$ ($\perp$) means implicit (explicit) rejection.

[2] The modular analysis in [HHK17] suggests that the FO implicitly contains the GEM/REACT at least the proof technique.

Quantum random oracle model

- FO-like generic constructions 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).
- Generic constructions in the **ROM** have gathered renewed interest in post-quantum setting, where adversaries are equipped with a quantum computer.
- In a real world, quantum adversary can execute hash functions (the instantiation of **RO**) on an arbitrary superposition of inputs.
- Therefore, as argued by Boneh et al. [BDF+11], when proving 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.

Quantum random oracle model

- In general, QROM is quite difficult to deal with, since many proof techniques in the ROM will be incompatible with the QROM.
- In the ROM, the simulator naturally "learns" the queries to the RO. This is called extractability, which is widely used in proving security for cryptosystem under computational hard problems in the indistinguishability security model.
- In the QROM, the queries can be quantum states, and "learning" a quantum state means a measurement, which allows to extract classical information from a quantum state.
- Separations of ROM and QROM were given by [BDF+11, YZ21].

Security reduction

When proving a security of a cryptographic scheme S under a hardness assumption of a problem P, we usually construct a reduction algorithm R against P that uses an adversary $\mathcal{A}$ against S as a subroutine.

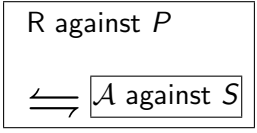

[3] This name comes from [GCS+17].

Security reduction

When proving a security of a cryptographic scheme S under a hardness assumption of a problem P, we usually construct a reduction algorithm R against P that uses an adversary $\mathcal{A}$ against S as a subroutine.

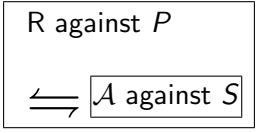

Query-based[3] The reduction uses a RO-query from the adversary to break the underlying hard problem.

[3] This name comes from [GCS+17].

Security reduction

When proving a security of a cryptographic scheme S under a hardness assumption of a problem P, we usually construct a reduction algorithm R against P that uses an adversary $\mathcal{A}$ against S as a subroutine.

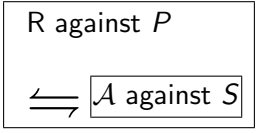

Query-based[3] The reduction uses a RO-query from the adversary to break the underlying hard problem.

Measurement-based The reduction measures a RO-query from the adversary and uses the measurement outcome to break the underlying hard problem.

Note: The measurement-based reduction is the quantum version of the query-based reduction.

[3] This name comes from [GCS+17].

Tightness

- Let $(T_{\mathcal{A}}, \epsilon_{\mathcal{A}})$ and (T_R, ϵ_R) denote the running times and advantages of $\mathcal{A}$ and R, respectively.
- The reduction is said to be *tight* if $T_{\mathcal{A}} \approx T_R$ and $\epsilon_{\mathcal{A}} \approx \epsilon_R$.
- Otherwise, if $T_R \gg T_{\mathcal{A}}$ or $\epsilon_R \ll \epsilon_{\mathcal{A}}$, the reduction is *non-tight*.
- The tightness gap, (informally) defined by $\frac{T_R \epsilon_{\mathcal{A}}}{T_{\mathcal{A}} \epsilon_R}$ [Men12], is used to measure the quality of a reduction.
- Tighter reductions with smaller tightness gap are desirable for practice cryptography especially in large-scale scenarios, since the tightness of a reduction determines the strength of the security guarantees provided by the security proof.

Black-box V.S. Non-black-box

Black-Box (BB) The reduction merely uses the adversary's input-output behavior, and does not depend on the internals

Non-Black-Box (NBB) The reduction requires knowledge of the adversary's internals like the adversary's code.

In general, black-box reductions are more pervasive than the non-black-box ones in cryptography.

Current proofs for FO-like KEM constructions

- Most QROM reductions (including black-box and non-black-box) for FO-like KEM constructions from standard CPA assumptions are measurement-based, and have the tightness[4]
  1. T_R is about $T_{\mathcal{A}}$;
  2. $\epsilon_R \approx \frac{1}{\kappa}\epsilon_{\mathcal{A}}^{\tau}$.

κ: the factor of security loss τ: the degree of security loss

[4]When comparing the tightness of different reductions, we assume perfect correctness of underlying scheme for brevity.

Current proofs for FO-like KEM constructions

Table: The tightness of current QROM proofs from standard CPA assumptions.

(κ, τ)	Variants of FO	Variants of REACT/GEM	Type
[HHK17]	$(q^6, 4)$[5]	$(q^2, 2)$	BB
[SXY18,**J**ZC+18]	$(q^2, 2)$	$(q^2, 2)$	BB
[**J**ZM19a,**J**ZM19b]	$(q, 2)$	$(q, 2)$	BB
[BHH19]	$(q, 2)$	$(1, 2)$	BB
[KSS+20]	$(q^2, 1)$	$(q, 1)$	NBB[6]

[5] q is the total number of adversary's queries (including quantum and classical) to various oracles.

[6] The reduction in [KSS+20] relies on a newly introduced "Measure-Rewind-Measure" (MRM) technique that can only apply to reversible adversaries. In post-quantum setting, most adversaries are irreversible since most oracles (e.g., decap. oracle) in the security model can only be classically queried.

Current proofs for FO-like KEM constructions

- As we can see, the existing black-box QROM reductions from standard CPA assumptions, are far from desirable due to the quadratic security loss (at least).
- Although this quadratic loss can be avoided by non-black-box reductions [KSS+20], the reductions in [KSS+20] can only apply to *reversible* adversaries.
- Note that the existing black-box reductions in the literature can cover arbitrary adversaries.
- These results are quite different from the ROM counterpart, where a linear loss can be achieved in a black-box manner [Den03, HHK17].

One-way to Hiding (OW2H)

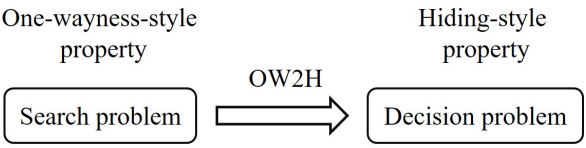

- The quadratic loss arises from the usage of the OW2H technique (an essential technique to prove post-quantum security) [Unr15].
- The OW2H technique has a quadratic loss. Very recently, several works [AHU18, BHH+19, KSS+20] tried to improve the tightness of OW2H.
- However, as in the case of FO-like KEMs, the tightness improvements are only restricted to the factor of reduction loss, and the quadratic loss still exists (except the non-black-box MRM-OW2H [KSS+20]).

Motivation

A natural question is that

For FO-like KEMs and the OW2H technique, is the quadratic loss unavoidable for measurement-based black-box reductions?

Main Contribution

In this paper, we give an affirmative answer for the above question.

- For FO-like KEMs, we show a measurement-based black-box reduction from breaking the standard OW-CPA (or IND-CPA) security of the underlying PKE to breaking the IND-CCA security of the resulting KEM, will *inevitably* incur a quadratic loss of the security.
- Such an impossibility result can also be extended to show that the quadratic loss is also unavoidable when one turns a search problem into a decision problem via the essential OW2H technique in a black-box manner. That is, the black-box OW2H technique [Unr15, AHU18, BHH+19] is essentially optimal in terms of the degree of reduction loss.

Main techniques

Here, we just take $\mathrm{KEM} - \mathrm{U}^{\not\perp}_{m}$ as an example. But it's not hard to extend the results to other FO-like KEM constructions and the general one-way to hiding.

Gen	$Encaps(pk)$	$Decaps(sk', c)$
1 : $(pk, sk) \leftarrow Gen'$	1 : $m \xleftarrow{\$} \mathcal{M}$	1 : Parse $sk' = (sk, k)$
2 : $k \xleftarrow{\$} \mathcal{K}^{prf}$	2 : $c := Enc'(pk, m)$	2 : $m' := Dec'(sk, c)$
3 : $sk' := (sk, k)$	3 : $K := H(m)$	3 : **if** $Enc'(pk, m') = c$
4 : **return** (pk, sk')	4 : **return** (K, c)	4 : **return** $K := H(m')$
		5 : **else return**
		6 : $K := f(k, c)$

Figure: IND-CCA-secure $\mathrm{KEM} - \mathrm{U}^{\not\perp}_{m} = \mathrm{U}^{\not\perp}_{m}[\mathrm{DPKE}, H, f]$

Proof skeleton

1. We first construct a specific quantum adversary $\mathcal{A}$ that breaks the IND-CCA security of the resulting KEM with advantage at least $\sqrt{p}$ (p is a real in $[0, 1]$), i.e., $\epsilon_{\mathcal{A}} \gtrapprox \sqrt{p}$.
2. Then, we show that any measurement-based black-box reduction $R^{\mathcal{A}}$ that runs this specific $\mathcal{A}$ as a subroutine to break the OW-CPA security of the underlying DPKE will have advantage at most p, i.e., $\epsilon_R \lessapprox p$.

Unbounded quantum adversary

- When attacking the IND-CCA security of $\mathrm{KEM} - \mathrm{U}^{\not\perp}_{m}$, an adversary $\mathcal{A}(pk, c^*, K_b)$ needs to distinguish $K_0 = H(m^*)$ from a uniformly random key K_1, where $c^* = Enc(pk, m^*)$ for a uniformly random m^*, the coin $b \in \{0, 1\}$ is uniformly random.
- We note that the random oracle H has a useful property that if m^* has not been queried to H by $\mathcal{A}$, then the value $H(m^*)$ is uniformly random in $\mathcal{A}$'s view. Thus, $\mathcal{A}$'s distinguishing advantage is negligible when making no queries to H with m^*.
- Intuitively, to achieve a non-negligible distinguishing advantage, $\mathcal{A}$ has to query m^* to H.

Unbounded quantum adversary

- In the ROM, $\mathcal{A}$ can only make classical queries to H.
- For any p ($0 \leq p \leq 1$), if $\mathcal{A}$ queries m^* to H with probability p, he will learn $K_0 = H(m^*)$ with probability p and break the IND-CCA security with advantage approximately p by testing whether $K_0 = K_b$.
- For a reduction $R^{\mathcal{A}}$ against the OW-CPA security of the underlying DPKE, a natural way is to take $\mathcal{A}$'s query as a return.
- Then, with probability p, $R^{\mathcal{A}}$ will return the m^* and break the OW-CPA security of the underlying DPKE.
- That is, the advantages of $R^{\mathcal{A}}$ and $\mathcal{A}$ are approximately equal, which is consistent with currently known tight reduction in [HHK17].

Unbounded quantum adversary

- In the QROM, a quantum adversary $\mathcal{A}$ can make a query to H with a quantum state. Consider the following quantum state

$$|\psi_{-1}\rangle := \sqrt{p}|m^*\rangle|0\rangle + \sqrt{1-p}|m'\rangle|\Sigma\rangle,$$

  where $m' \neq m^*$, $|\Sigma\rangle = \sum_{k\in\mathcal{K}} \frac{1}{\sqrt{|\mathcal{K}|}}|k\rangle$ and $\mathcal{K}$ is the (session) key space.

- For a quantum query with $|\psi_{-1}\rangle$, the random oracle H will return

$$|\psi_0\rangle := \sqrt{p}|m^*\rangle|K_0\rangle + \sqrt{1-p}|m'\rangle|\Sigma\rangle.$$

Remark: If the adversary $\mathcal{A}$ directly measures $|\psi_0\rangle$ in standard computational basis, he will obtain K_0 with probability p and break the IND-CCA security with the advantage (approximately) p by testing whether $K_0 = K_b$ as the adversary in the ROM.

Unbounded quantum adversary

- A quantum adversary $\mathcal{A}$ can directly guess b by *testing* whether the quantum state $|\psi_0\rangle$ is equal to quantum state $|\psi_b\rangle$, where

$$|\psi_b\rangle := \sqrt{p}|m^*\rangle|K_b\rangle + \sqrt{1-p}|m'\rangle|\Sigma\rangle.$$

- Testing whether $|\psi_0\rangle$ is equal to $|\psi_b\rangle$[7] can be accomplished using the standard quantum state discrimination method (known as Helstrom measurement) [Hel79] with advantage (approximately) at least $\sqrt{p}$.

[7] Formally, we need to judge $|\psi_0\rangle\langle\psi_0|$ comes from $|\psi_b\rangle\langle\psi_b|$ or $\mathbb{E}_{K_{1-b}}|\psi_{1-b}\rangle\langle\psi_{1-b}|$ (the the expectation is taken over $K_{1-b} \xleftarrow{\$} \mathcal{K}$).

Unbounded quantum adversary

The unbounded quantum adversary $\mathcal{A}(pk, c^*, K_b)$ is as follows.

$\mathcal{A}(pk, c^*, K_b)$

1 : Search a $m^* \in \mathcal{M}$ such that $Enc'(pk, m^*) = c^*$

// If no one (or more than one) is found, output 1 and terminate the procedure.

2 : Sample a real $p \in [0, 1]$ and a uniform m' from $\{m' \in \mathcal{M} : m' \neq m^*\}$

3 : Query H with quantum state $|\psi_{-1}\rangle := \sqrt{p}|m^*\rangle|0\rangle + \sqrt{1-p}|m'\rangle|\Sigma\rangle$

4 : Perform Helstrom measurement M on $|\psi_0\rangle$ (the state returned by H)

5 : Return the measurement outcome.

Theorem 3.1 (The advantage of $\mathcal{A}$ in the QROM).

If the underlying DPKE is perfectly correct, the advantage of $\mathcal{A}$ against the IND-CCA security of $\mathrm{KEM} - \mathrm{U}_m^{\not\perp}$ is at least $\sqrt{p}(1 - 1/|\mathcal{K}|) \approx \sqrt{p}$.

The advantage of a measurement-based reduction

Measurement-based (black-box) reduction

1. Reduction R receives a challenge $inpt_1$ as input, runs a PPT preprocessing (quantum) subalgorithm $(inpt, rand, s) \leftarrow R_1(inpt_1)$, and then launches $\mathcal{A}(inpt; rand)$[8].
2. When $\mathcal{A}$ makes a query to the RO with quantum state ϕ, R measures ϕ in the computational basis[9], and gets the measurement outcome $mest$.
3. Reduction R runs a PPT postprocessing (quantum) subalgorithm $out \leftarrow R_2(s, mest)$, and returns out.

[8] Here, $inpt_1$, $inpt$ and $rand$ are classical, and s can be a quantum state.

[9] The reduction R just measures the query input registers.

The advantage of a measurement-based black-box reduction

Measurement-based (black-box) reduction

Remark 1.

Performing an additional quantum (unitary) operation on adversary's query before measuring isn't allowed. But, such an additional unitary operation cannot substantially increase reduction's advantage, otherwise there exists an algorithm breaking the OW-CPA security of the underlying DPKE efficiently.

Remark 2.

Note that the considered reductions do not restrict the simulations of random oracles and other oracles that adversary queries, and thus can cover the black-box reductions in [HHK17, SXY18, JZC+18, JZM19a, JZM19b, BHH+19, HKSU20].

The advantage of a measurement-based reduction

Meta-reduction methodology

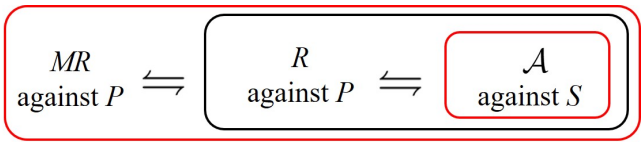

- Meta-reduction methodology has proven to be a versatile tool in deriving impossibility results and tightness bounds of security proofs.
- A meta-reduction MR_R simulates the adversarial part $\mathcal{A}$, runs the reduction R as a subroutine, and break the underlying hard problem P directly.
- That is, a meta-reduction MR_R treats the reduction R as an adversary itself and reduce the existence of such a reduction R to a presumably hard problem.

The advantage of a measurement-based reduction

Consider the advantage of $R^{\mathcal{A}}$ in following three cases, where $\textsc{Ine}$ ($\textsc{Exi}$, resp.) is the event that the exhaustive search returns no (a, resp.) m^* such that $Enc(pk, m^*) = c^*$, and $\textsc{Good}$ ($\textsc{Bad}$, resp.) is the event that the measurement outcome is (not, resp.) m^*.

Case 1: $\textsc{Ine}$. In this case, $\mathcal{A}$ just outputs 1 without queries to H. Thus, exhaustive search for m^* in this case is vain, and $\mathcal{A}$ can be replaced by an adversary $\mathcal{A}_1$ that always outputs 1 without the search for m^* and the query to the random oracle H. Therefore, we can easily construct a meta-reduction MR_1^R that simulates $\mathcal{A}_1$ and takes $R^{\mathcal{A}_1}$ as a subroutine to break the OW-CPA security of the underlying DPKE such that the running time of MR_1^R is about the running time of R, and under the condition $\textsc{Ine}$ the advantage of MR_1^R is about the advantage of R.

Case 2: $\textsc{Exi} \wedge \textsc{Good}$. Since $\Pr[\textsc{Good}|\textsc{Exi}] = p$, we can bound the advantage of R in this case by p.

The advantage of a measurement-based reduction

Consider the advantage of $R^{\mathcal{A}}$ in following three cases, where $\textsc{Ine}$ ($\textsc{Exi}$, resp.) is the event that the exhaustive search returns no (a, resp.) m^* such that $Enc(pk, m^*) = c^*$, and $\textsc{Good}$ ($\textsc{Bad}$, resp.) is the event that the measurement outcome is (not, resp.) m^*.

Case 3: $\textsc{Exi} \wedge \textsc{Bad}$. In this case, R gets $m' \neq m^*$. Let $\mathcal{A}_2$ be an adversary that queries a quantum state $\sum_{m,k} \frac{1}{\sqrt{|\mathcal{M}|\cdot|\mathcal{K}|}}|m\rangle|k\rangle$ and outputs 1 without the search for m^*. Thus, the advantage of R under the condition $\textsc{Exi} \wedge \textsc{Bad}$ remains unchanged when $\mathcal{A}$ is replaced by $\mathcal{A}_2$. As in the case 1, we can also construct a meta-reduction MR_2^R against the OW-CPA security of the underlying DPKE that simulates $\mathcal{A}_2$ and takes $R^{\mathcal{A}_2}$ as a subroutine such that the running time of MR_2^R is about the running time of R, and under the condition $\textsc{Exi} \wedge \textsc{Bad}$ the advantage of MR_2^R is about the advantage of R.

The advantage of a measurement-based reduction

Theorem 3.2.

If the underlying DPKE is perfectly correct, for any measurement-based reduction $R^{\mathcal{A}}$ that runs the adversary $\mathcal{A}$ once without rewinding, there exist two meta-reductions MR_1^R and MR_2^R against the OW-CPA security of the underlying DPKE such that

$$\epsilon_R \leq p + \epsilon_{MR_1} + \frac{|\mathcal{M}|}{|\mathcal{M}| - 1}\epsilon_{MR_2},$$

and $Time(R) \approx Time(MR_1) \approx Time(MR_2)$.

Main theorem

Combing Theorems 3.1 and 3.2, we can directly obtain the following main Theorem.

Theorem 3.3.

If the underlying DPKE is perfectly correct, there exists a quantum adversary $\mathcal{A}$ against the IND-CCA security of $\mathrm{KEM} - \mathrm{U}_m^{\not\perp}$ such that for any measurement-based (black-box) reduction $R^{\mathcal{A}}$ that runs $\mathcal{A}$ (once without rewinding), measures $\mathcal{A}$'s query and uses the measurement outcome to break the OW-CPA security of the underlying DPKE, there exist two meta-reductions MR_1^R and MR_2^R which take R as a subroutine to break the OW-CPA security of the underlying DPKE such that

$$\epsilon_{\mathcal{A}} \geq (1 - \frac{1}{|\mathcal{K}|}) \times \sqrt{\epsilon_R - \epsilon_{MR_1} - \frac{|\mathcal{M}|}{|\mathcal{M}| - 1} \cdot \epsilon_{MR_2}}$$

and $Time(R) \approx Time(MR_1^R) \approx Time(MR_2^R)$.

Interpretation

- Under the assumption that the advantage of any efficient algorithm breaking the OW-CPA security of the underlying DPKE is negligible, we have ϵ_{MR_1} and ϵ_{MR_2} are negligible.

Interpretation

- Under the assumption that the advantage of any efficient algorithm breaking the OW-CPA security of the underlying DPKE is negligible, we have ϵ_{MR_1} and ϵ_{MR_2} are negligible.

Thus, we have

$$\epsilon_R \lessapprox \epsilon_{\mathcal{A}}^2.$$

For $\mathrm{KEM} - \mathrm{U}_m^{\not\perp}$, a measurement-based black-box reduction in the QROM from breaking standard OW-CPA security of the underlying DPKE to breaking the IND-CCA security of the resulting KEM, will *inevitably* incur a quadratic loss of the security.

Conclusion

- For FO-like KEMs, we first show the tightness limits of the black-box reductions, and prove that a *measurement-based* reduction in the QROM from breaking the standard CPA security of the underlying PKE to breaking the IND-CCA security of the resulting KEM, will *inevitably* incur a quadratic loss of the security.
- In particular, most black-box reductions for these FO-like KEMs are of this type, and our results suggest an explanation for the lack of progress in improving this reduction tightness in terms of the degree of security loss.
- This impossibility results can also be extended to show the tightness limits of the general (black-box) one-way to hiding.

Thanks for your attention!

hdjiang13@gmail.com

References

AHU18 Andris Ambainis, Mike Hamburg and Dominique Unruh, Quantum security proofs using semi-classical oracles

ARU14 Andris Ambainis, Ansis Rosmanis and Dominique Unruh, Quantum attacks on classical proof systems: The hardness of quantum rewinding

BDF+11 Dan Boneh et al., Random oracles in a quantum world

BHH+19 Nina Bindel, Mike Hamburg, Kathrin Hövelmanns, Andreas Hülsing, and Edoardo Persichetti, Tighter proofs of CCA security in the quantum random oracle model

Den03 Alexander W. Dent, A designers guide to KEMs

GCS+17 Fuchun Guo et al., Optimal security reductions for unique signatures: bypassing impossibilities with a counterexample

References

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

HKSU20 Kathrin Hövelmanns, Eike Kiltz, Sven Schäge and Dominique Unruh, Generic Authenticated Key Exchange in the Quantum Random Oracle Model

Hel79 Carl W. Helstrom, Quantum detection and estimation theory

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

JZM19a Haodong Jiang, Zhenfeng Zhang and Zhi Ma. Key encapsulation mechanism with explicit rejection in the quantum random oracle model

JZM19b Haodong Jiang, Zhenfeng Zhang and Zhi Ma. Tighter security proofs for generic key encapsulation mechanism in the quantum random oracle model

References

KSS+20 Measure-rewind-measure: tighter quantum random oracle model proofs for one-way to hiding lemma and CCA security

Men12 Menezes Alfred, Another Look at Provable Security

Unr15 Dominique Unruh, Non-Interactive Zero-Knowledge Proofs in the Quantum Random Oracle Model

SXY18 Tsunekazu Saito, Keita Xagawa and Takashi Yamakawa, Tightly-secure key-encapsulation mechanism in the quantum random oracle model

YZ21 Takashi Yamakawa and Mark Zhandry, Classical vs quantum random oracles

Zha19 Mark Zhandry, How to record quantum queries, and applications to quantum indifferentiability