CCA encryption in the QROM, pt. I

Known security statements for CCA transformations

Kathrin Hövelmanns[1]

[1]Horst Görtz Institute for IT Security, Ruhr University Bochum, Germany

Berkeley lattice workshop: From Theory to Practice
April 28th, 2020

Context: NIST 'competition'

Goal: Quantum-secure public-key encryption and signatures

Desired: Active security (CCA)

Easier to achieve: Passive security (OW/CPA)

Can we turn passive into active, generically?

Frequently used solution: FO transformation [FO99,13] and its variants

Originally proven in random oracle model

This talk: What happens if quantum adversary interacts with (non-quantum) network?

Outline

Goal of this talk: Preparation for next talk

$\rightarrow$ No newness, but a survey:

1. Reminder: Quantum ROM and Oneway-to-Hiding (OWTH)
2. Overview: FO-like transformations and known security results
    - Results for deterministic schemes
    - Results with derandomisation
3. Does OWTH imply quadratic loss?

Security reductions and (quantum) Random Oracles

Random Oracle Model (ROM)

Proof heuristic: Replace hash fct. with perfectly random fct. H

Common proof strategy:

A can distinguish $\mathsf{H}(x^*)$ from random

$\Rightarrow$ Reduction learns preimage x^* (and x^* solves P)

What if A is quantum?

Quantum Random Oracle Model (QROM) [BDFLSZ10]

Scenario: Quantum adversary interacting with non-quantum network $\Rightarrow$

- "Online" primitives (decryption, signing, ...) stay classical
- "Offline" primitives (like hash functions) computable in superposition

What's new: A might evaluate hash function on some superposition

$$\sum_{x \in X} \alpha_x |x\rangle$$

Superposition: Function's domain X gives rise to vector space $\mathbb{C}^X$

Quantum state = Linear combination of base vectors $|x\rangle$ s. th.

$$\sum_{x \in X} |\alpha_x|^2 = 1$$

How do we formalise quantum-accessibility of the random oracle?

Quantum Random Oracle Model (QROM) [BDFLSZ10]

Model quantum-accessible version of O by mapping U_{O}:

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

where x (y) are base states of the input (output) register

Model $\mathrm{A}^{|\mathrm{O}\rangle}$ via sequence of attack unitaries A_i, interleaved with oracle queries:

$$\mathrm{A}^{|\mathrm{O}\rangle} \mathrel{\widehat{=}} \mathrm{A}_N \circ U_{\mathrm{O}} \circ \mathrm{A}_{N-1} \circ \cdots \circ U_{\mathrm{O}} \circ \mathrm{A}_1$$

(ith random oracle query $\mathrel{\widehat{=}}$ output of A_i)

Question: How to extract a particular preimage from a query?

Original "Oneway to Hiding" [Unruh14]

Quantum generalisation of "random-until-QUERY":

$$\left|\Pr\left[1 \leftarrow \mathsf{A}^{|\mathsf{O}\rangle}(x^*, \mathsf{O}(x^*))\right] - \Pr\left[1 \leftarrow \mathsf{A}^{|\mathsf{O}\rangle}(x^*, \$)\right]\right| \leq 2q \cdot \sqrt{\epsilon}$$

where

$$\epsilon := \Pr\left[\text{Measuring a random query gives us } x^*\right]$$

Tightness improvements for OWTH:

Variant	Bound	Additional restrictions
Original (above)	$2q\sqrt{\epsilon}$	
Semi-classical [AHU18]	$2\sqrt{q\epsilon}$	✓
Double-sided [BH+19]	$2\sqrt{\epsilon}$	✓
Next talk [KS+20]	$4q\epsilon$	✓

Overview: FO-like transformations and current results

Common ground of all recent modularisations

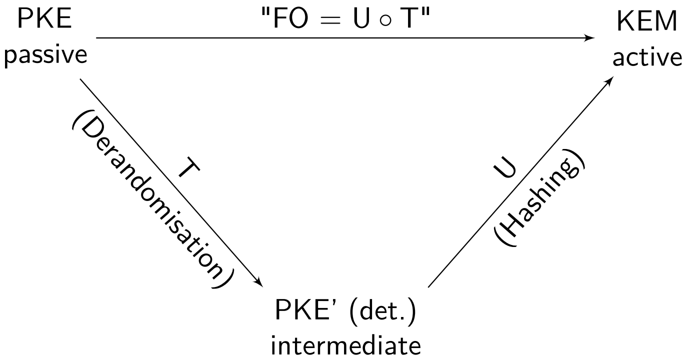

At least one step uses OWTH

Many different variations of U

$U^{\perp}$

- Encapsulation:
  1. Choose uniformly random plaintext m
  2. Use Enc' to encrypt m to ciphertext c
  3. $k := \mathsf{H}(m, c)$

Many different variations of U

$U^{\perp}$

- Encapsulation:
  1. Choose uniformly random plaintext m
  2. Use Enc' to encrypt m to ciphertext c
  3. $k := \mathrm{H}(m, c)$

- Decapsulation:
  1. Use Dec' to decrypt c to plaintext m'
  2. If c decrypts to $\perp$
  3. return $\perp$
  4. return $k' := \mathrm{H}(m', c)$

Many different variations of U

$\mathsf{U}_m^{\perp}$

- Encapsulation:
  1. Choose uniformly random plaintext m
  2. Use Enc' to encrypt m to ciphertext c
  3. $k := \mathsf{H}(m, c)$ $\mathsf{H}(m)$

- Decapsulation:
  1. Use Dec' to decrypt c to plaintext m'
  2. If c decrypts to $\perp$
  3. return $\perp$
  4. return $k' := \mathsf{H}(m', c)$ $\mathsf{H}(m')$

Many different variations of U

$\mathsf{U}_{m}^{\not\perp}$

- Encapsulation:
  1. Choose uniformly random plaintext m
  2. Use Enc' to encrypt m to ciphertext c
  3. $k := \mathsf{H}(m, c)$ $\mathsf{H}(m)$

- Decapsulation:
  1. Use Dec' to decrypt c to plaintext m'
  2. If c decrypts to $\perp$
  3. return $\perp$ return pseudorandom value ("implicit rejection")
  4. return $k' := \mathsf{H}(m', c)$ $\mathsf{H}(m')$

Many different variations of U

$U_{m}^{\not\perp \circlearrowleft}$

- Encapsulation:
  1. Choose uniformly random plaintext m
  2. Use Enc' to encrypt m to ciphertext c
  3. $k := \mathrm{H}(m, c)$ $\mathrm{H}(m)$

- Decapsulation:
  1. Use Dec' to decrypt c to plaintext m'
  2. If c decrypts to $\perp$ or $\mathrm{Enc}'(m') \neq c$ ("reencryption")
  3. return $\perp$ return pseudorandom value ("implicit rejection")
  4. return $k' := \mathrm{H}(m', c)$ $\mathrm{H}(m')$

Cave: New reencryption step not always emphasised!

Many different variations of U

$\mathsf{U}_{m}^{\not\perp\circlearrowleft}$-KC

- Encapsulation:
  1. Choose uniformly random plaintext m
  2. Use Enc$'$ to encrypt m to ciphertext c
  3. $k := \mathsf{H}(m, c)$ $\mathsf{H}(m)$
  4. Append to c a "key confirmation ciphertext" $d := \mathsf{H}'(m)$

- Decapsulation:
  1. Use Dec$'$ to decrypt c to plaintext m'
  2. If c decrypts to $\perp$ or $\mathsf{Enc}'(m') \neq c$ or $\mathsf{H}'(m') \neq d$
  3. return $\perp$ return pseudorandom value ("implicit rejection")
  4. return $k' := \mathsf{H}(m', c)$ $\mathsf{H}(m')$

Cave: New reencryption step not always emphasised!

Common ground of all recent modularisations

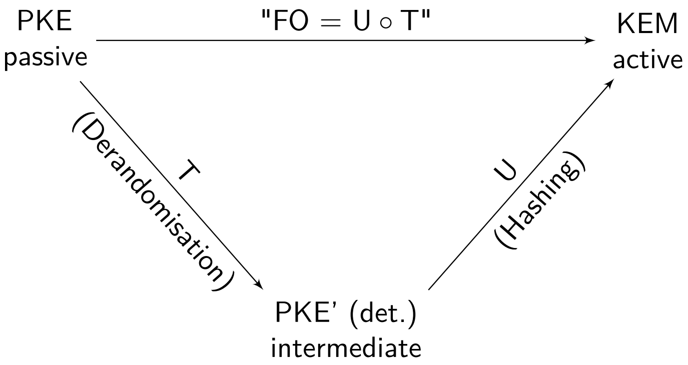

At least one step uses OWTH

Deterministic schemes

Applying U to deterministic schemes: State of the art

SXY18: PKE$'$ perf. correct and disjoint simulatable $\rightarrow$ tight CCA security

Disjoint simulatability: Efficiently sampleable "fake ciphertexts" s.th.

1. fake cts indistinguishable from real cts
2. fake cts invalid w.o.p

Applying U to deterministic schemes: State of the art

SXY18: PKE' perf. correct and disjoint simulatable $\rightarrow$ tight CCA security

Disjoint simulatability: Efficiently sampleable “fake ciphertexts” s.th.

1. fake cts indistinguishable from real cts
2. fake cts invalid w.o.p

Intuition: Disjoint simulatability $\rightarrow$ can circumvent OWTH

perfect correctness required for consistency

generalisation not straightforward ☹

Applying U to deterministic schemes: State of the art

PKE$'$ FFC and η-injective $\rightarrow$ CCA security with
quadratic loss in the advantage [BHHHP19] or
linear loss in the number of RO queries [KS+20] (next talk)

FFC: Hard to find a valid ciphertext that decrypts incorrectly

η-injective: Enc$'$ is injective w.p. $1 - \eta$

Applying U to deterministic schemes: State of the art

All results use reencryption (= use $\mathsf{U}^{\circlearrowleft}$-variant)

Equivalency for implicit reject ($\mathsf{U}^{\not\perp}$): We can derive the key

via $k = \mathsf{H}(m, c)$ (= use $\mathsf{U}^{\not\perp,\circlearrowleft}$)

via $k = \mathsf{H}(m)$ (= use $\mathsf{U}_m^{\not\perp,\circlearrowleft}$)

Implication for explicit reject ($\mathsf{U}^{\perp}$):

Works for U_m-variant if we add key confirmation (= use $\mathsf{U}_m^{\perp\circlearrowleft}$-KC)

Applying U to deterministic schemes: Proof overview

Variant	Notion	Correctness	Add. requ.	CCA Bound (simplified)	How
$\mathsf{U}_m^{\not\perp\circlearrowleft}$	DS (det.)	perfect		tight	SXY18, Th. 4.2
$\mathsf{U}_m^{\perp\circlearrowleft}$-KC	DS (det.)	perfect		tight	JZM19a, Th. 5

Applying U to deterministic schemes: Proof overview

Variant	Notion	Correctness	Add. requ.	CCA Bound (simplified)	How
$\mathsf{U}_m^{\not\perp\circlearrowleft}$	DS (det.)	perfect		tight	SXY18, Th. 4.2
$\mathsf{U}_m^{\perp\circlearrowleft}$-KC	DS (det.)	perfect		tight	JZM19a, Th. 5
$\mathsf{U}^{\not\perp\circlearrowleft}$	OW (det.)	FFC	η-inj.	$\sqrt{\mathsf{OW}}$	BH+19, Th. 2
				or $q \cdot \mathsf{OW}$	KS+20 (next talk)

Tradeoff: generality vs tightness

Applying U to deterministic schemes: Proof overview

Variant	Notion	Correctness	Add. requ.	CCA Bound (simplified)	How
$\mathsf{U}_m^{\not\perp\circlearrowleft}$	DS (det.)	perfect		tight	SXY18, Th. 4.2
$\mathsf{U}_m^{\perp\circlearrowleft}$-KC	DS (det.)	perfect		tight	JZM19a, Th. 5
$\mathsf{U}^{\not\perp\circlearrowleft}$	OW (det.)	FFC	η-inj.	$\sqrt{\mathrm{OW}}$	BH+19, Th. 2
				or $q \cdot \mathrm{OW}$	KS+20 (next talk)

Tradeoff: generality vs tightness

Applying [BH+19, Ths. 5 and 4] leads to the following corollaries:

Variant	Notion	Correctness	Add. requ.	CCA Bound (simplified)	How
$\mathsf{U}^{\not\perp\circlearrowleft}$	DS (det.)	perfect		tight	Th. 5
$\mathsf{U}_m^{\not\perp\circlearrowleft}$	OW (det.)	FFC	η-inj.	$\sqrt{\mathrm{OW}}$, $q \cdot \mathrm{OW}$	Th. 5
$\mathsf{U}_m^{\perp\circlearrowleft}$-KC	OW (det.)	FFC	η-inj.	$\sqrt{\mathrm{OW}}$, $q \cdot \mathrm{OW}$	Th. 5, then 4

Derandomisation

Common ground of all recent modularisations

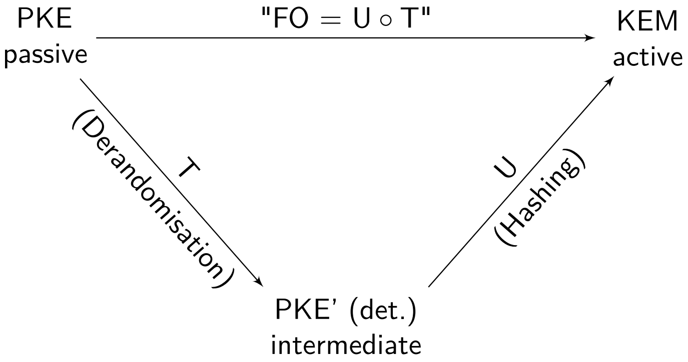

At least one step uses OWTH

Applying FO variants: State of the art

Diverse variants (like U-variants)

Recent tightness improvements for U $\Rightarrow$ Improvements for FO

Even nonmodular proofs imply security of other variants ([BH+19])

Applying FO variants: Proof overview

All results work for δ-correctness, require sufficiently large $\mathcal{M}$

Variant	Notion	Add. requ.	CCA Bound (simplified)	How
$\mathrm{FO}^{\not\perp}_{(m)}$	OW		$q\sqrt{\mathrm{OW}} + q\sqrt{\delta}$	JZ+18, Ths. 1, 2
$\mathrm{FO}^{\perp}_{(m)}$-KC				JZM19a, Ths. 2, 4
$\mathrm{FO}^{\perp}_{(m)}$-KC	CPA		$\sqrt{q \cdot \mathrm{CPA}} + q\sqrt{\delta}$	JZM19a, Ths. 1, 3

Applying FO variants: Proof overview

All results work for δ-correctness, require sufficiently large $\mathcal{M}$

Variant	Notion	Add. requ.	CCA Bound (simplified)	How
$\mathsf{FO}^{\not\perp}_{(m)}$	OW		$q\sqrt{\mathsf{OW}} + q\sqrt{\delta}$	JZ+18, Ths. 1, 2
$\mathsf{FO}^{\perp}_{(m)}$-KC				JZM19a, Ths. 2, 4
$\mathsf{FO}^{\perp}_{(m)}$-KC	CPA		$\sqrt{q \cdot \mathsf{CPA}} + q\sqrt{\delta}$	JZM19a, Ths. 1, 3
$\mathsf{FO}^{\not\perp}_{m}$	CPA	DS	$\sqrt{q \cdot \mathsf{CPA}} + \mathsf{DS} + q^2\delta$	HK+18, Th. 3.2
	CPA	Punct.	$\sqrt{q \cdot \mathsf{CPA}} + q^2\delta$	HK+18, Th. 3.6

DS: ciphertexts (disjoint) simulatable

Puncturing: Removing one message from $\mathcal{M}$ achieves DS, generically

Applying FO variants: Proof overview

All results work for δ-correctness, require sufficiently large $\mathcal{M}$

Variant	Notion	Add. requ.	CCA Bound (simplified)	How
$\mathsf{FO}^{\not\perp}_{(m)}$	OW		$q\sqrt{\mathsf{OW}} + q\sqrt{\delta}$	JZ+18, Ths. 1, 2
$\mathsf{FO}^{\perp}_{(m)}$-KC				JZM19a, Ths. 2, 4
$\mathsf{FO}^{\perp}_{(m)}$-KC	CPA		$\sqrt{q \cdot \mathsf{CPA}} + q\sqrt{\delta}$	JZM19a, Ths. 1, 3
$\mathsf{FO}^{\not\perp}_{(m)}$, $\mathsf{FO}^{\perp}_{m}$-KC	CPA	DS	$\sqrt{q \cdot \mathsf{CPA}} + \mathsf{DS} + q^2\delta$	HK+18, Th. 3.2
	CPA	Punct.	$\sqrt{q \cdot \mathsf{CPA}} + q^2\delta$	HK+18, Th. 3.6

DS: ciphertexts (disjoint) simulatable

Puncturing: Removing one message from $\mathcal{M}$ achieves DS, generically

(These results are derived via BH+19, Ths. 4 and 5)

Applying FO variants: Proof overview

All results work for δ-correctness, require sufficiently large $\mathcal{M}$

Variant	Notion	Add. requ.	CCA Bound (simplified)	How
$\mathrm{FO}^{\not\perp}_{(m)}$	OW		$q\sqrt{\mathrm{OW}} + q\sqrt{\delta}$	JZ+18, Ths. 1, 2
$\mathrm{FO}^{\perp}_{(m)}$-KC				JZM19a, Ths. 2, 4
$\mathrm{FO}^{\perp}_{(m)}$-KC	CPA		$\sqrt{q \cdot \mathrm{CPA}} + q\sqrt{\delta}$	JZM19a, Ths. 1, 3
$\mathrm{FO}^{\not\perp}_{(m)}$, $\mathrm{FO}^{\perp}_{m}$-KC	CPA	DS	$\sqrt{q \cdot \mathrm{CPA}} + \mathrm{DS} + q^2\delta$	HK+18, Th. 3.2
	CPA	Punct.	$\sqrt{q \cdot \mathrm{CPA}} + q^2\delta$	HK+18, Th. 3.6
$\mathrm{FO}^{\not\perp}_{(m)}$, $\mathrm{FO}^{\perp}_{m}$-KC	CPA	INJ	$\sqrt{q \cdot \mathrm{CPA}} + q^2\delta$	BH+19, Ths. 1 + 2 + Lm. 6
			or $q^2 \cdot \mathrm{CPA} + q^2\delta$	replace Th. 2 with next talk

DS: ciphertexts (disjoint) simulatable

Puncturing: Removing one message from $\mathcal{M}$ achieves DS, generically

INJ : T[PKE] is η-injective

(These results are derived via BH+19, Ths. 4 and 5)

Does OWTH imply quadratic loss?

Last year's impossibility result [JZM19b]

One of the '10 questions': Is the sqare root meaningful?

BH+19: It might be impossible to avoid [JZM19b]

Apparently, it is not! (next talk)

So, how do we place the result of [JZM19b]?

Last year's impossibility result [JZM19b]

Reminder: $\mathsf{A}^{|\mathsf{O}\rangle}$ modeled via

$$\mathsf{A}_N \circ U_\mathsf{O} \circ \mathsf{A}_{N-1} \circ \cdots \circ U_\mathsf{O} \circ \mathsf{A}_1$$

(ith random oracle query $\widehat{=}$ output of A_i)

All OWTH applications until [KS+20]:

Extract preimage from oracle queries $\widehat{=}$ output register of A_i

$\rightarrow$ only considers input/output behaviour of A

[JZM19b]: This 'query extraction' approach leads to quadratic loss

New approach: Also consider A's internal workings:

A has to measure to recognise the difference between $\mathsf{O}(x^*)$ and \$
$\rightarrow$ Measurement reveals x^*

References

SXY18: Tightly-Secure Key-Encapsulation Mechanism in the Quantum Random Oracle Model, eprint: 2017/1005

JZCMW18: IND-CCA-secure Key Encapsulation Mechanism in the Quantum Random Oracle Model, Revisited, eprint: 2017/1096

HKSU18: Generic Authenticated Key Exchange in the Quantum Random Oracle Model, eprint: 2018/928

JZM19a: Key Encapsulation Mechanism with Explicit Rejection in the Quantum Random Oracle Model, eprint: 2019/052

JZM19b: On the non-tightness of measurement-based reductions for key encapsulation mechanism in the quantum random oracle model, eprint: 2019/494

BH+19: Tighter proofs of CCA security in the quantum random oracle model, eprint: 2019/590

KS+20: Measure-Rewind-Measure: Tighter Quantum Random Oracle Model Proofs for One-Way to Hiding and CCA Security (to appear)