Asiacrypt, Taipei, Taiwan

9/12/2022

Failing gracefully: Decryption failures and the Fujisaki-Okamoto transform

Kathrin Hövelmanns Andreas Hülsing Christian Majenz

Technical University of Denmark

EINDHOVEN UNIVERSITY OF TECHNOLOGY

Motivation

Computational problem
(LWE, NTRU, SD)…

PKE
Passively secure
(OW/IND-CPA)

Key Encapsulation
IND-CCA

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

Motivation

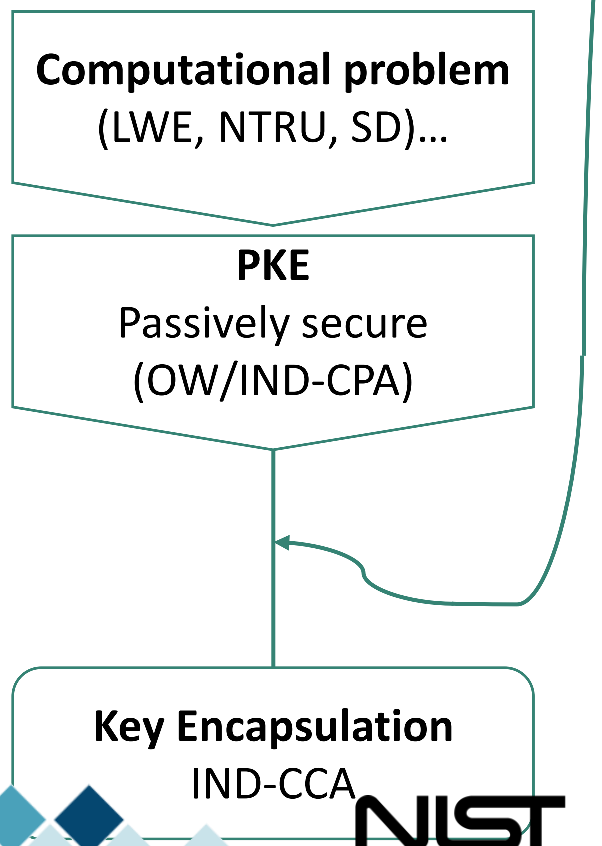

Fujisaki-Okamoto transform

Originally (FO99): no decryption failures (lattices, codes ☹)

Revisited (HHK17):

☑ small failure probability δ

different rejection methods

National Institute of Standards and Technology
U.S. Department of Commerce

Motivation

Computational problem
(LWE, NTRU, SD)…

PKE
Passively secure
(OW/IND-CPA)

Key Encapsulation
IND-CCA

Fujisaki-Okamoto transform

Originally (FO99): no decryption failures (lattices, codes ☹)

Revisited (HHK17):
☑ small failure probability δ
different rejection methods

Weird QROM thing 1

ROM: Rejection-method-agnostic

Quantum ROM:
Different methods → bounds vastly differ

National Institute of Standards and Technology
U.S. Department of Commerce

Motivation

Computational problem
(LWE, NTRU, SD)…

PKE
Passively secure
(OW/IND-CPA)

Key Encapsulation
IND-CCA

Fujisaki-Okamoto transform

Originally (FO99): no decryption failures (lattices, codes ☹)

Revisited (HHK17):
☑ small failure probability δ
different rejection methods

Weird QROM thing 1

ROM: Rejection-method-agnostic

Quantum ROM:
Different methods → bounds vastly differ

Weird QROM thing 2

Grover-like δ – term: $q^2 \cdot \delta$

…can attackers quantum search?

Suboptimal bounds?

National Institute of Standards and Technology
U.S. Department of Commerce

Motivation

Computational problem
(LWE, NTRU, SD)…

PKE
Passively secure
(OW/IND-CPA)

Key Encapsulation
IND-CCA

Fujisaki-Okamoto transform

Originally (FO99): no decryption failures (lattices, codes ☹)

Revisited (HHK17):
☑ small failure probability δ
different rejection methods

Weird QROM thing 1

ROM: Rejection-method-agnostic

Quantum ROM:
Different methods → bounds vastly differ

Weird QROM thing 2

Grover-like δ – term: $q^2 \cdot \delta$

…can attackers quantum search?

Suboptimal bounds?

Applicability issue

Concrete δ – estimations ⚡ security proofs

K Y B E R

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

δ - estimations vs security proofs

$\delta \triangleq$ advantage in

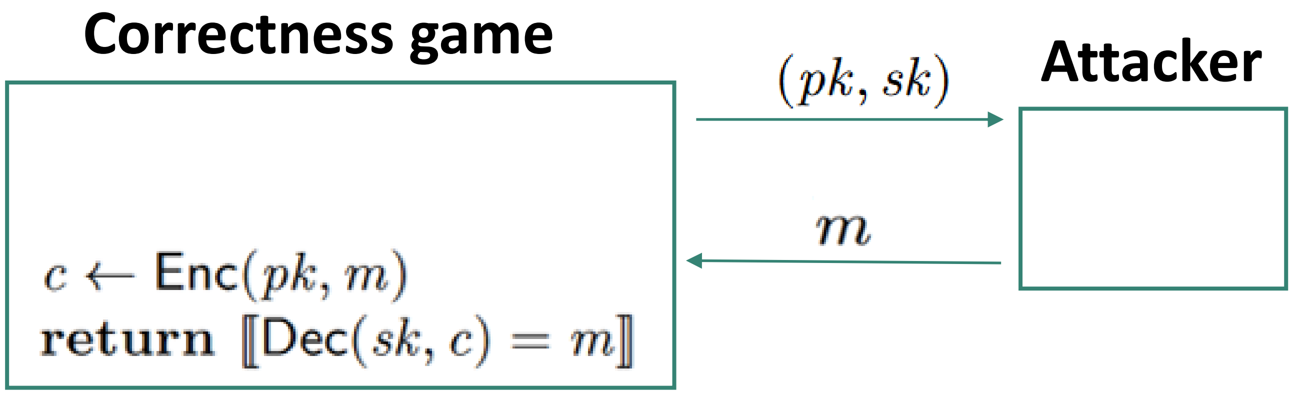

Applicability issue

Concrete δ – estimations ↯ security proofs

δ - estimations vs security proofs

$\delta \triangleq$ advantage in

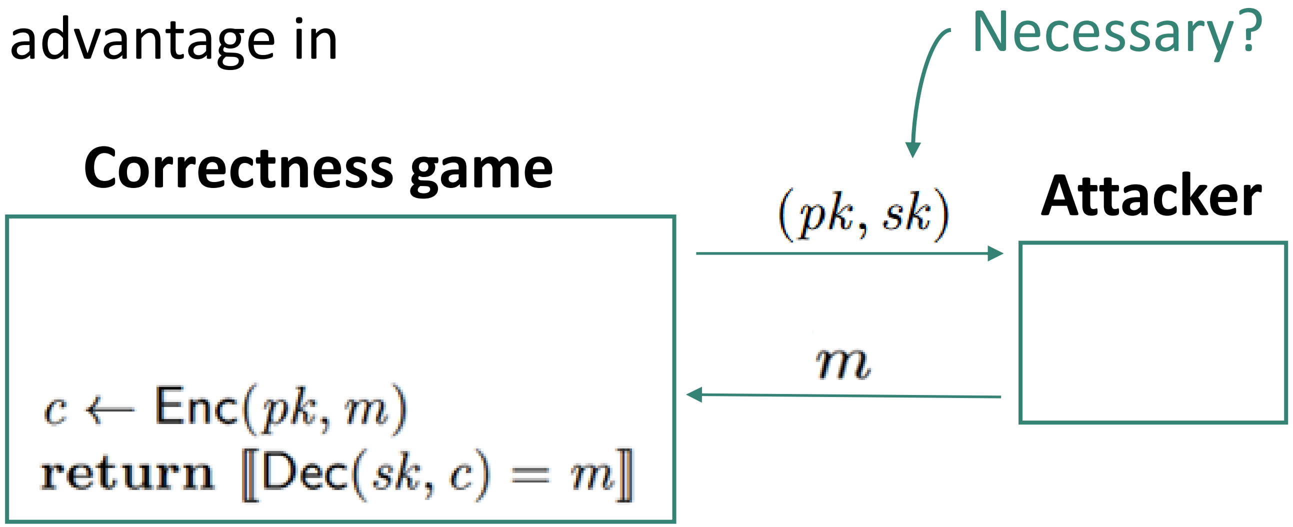

Applicability issue

Concrete δ – estimations ⚡ security proofs

δ - estimations vs security proofs

$\delta \triangleq$ advantage in

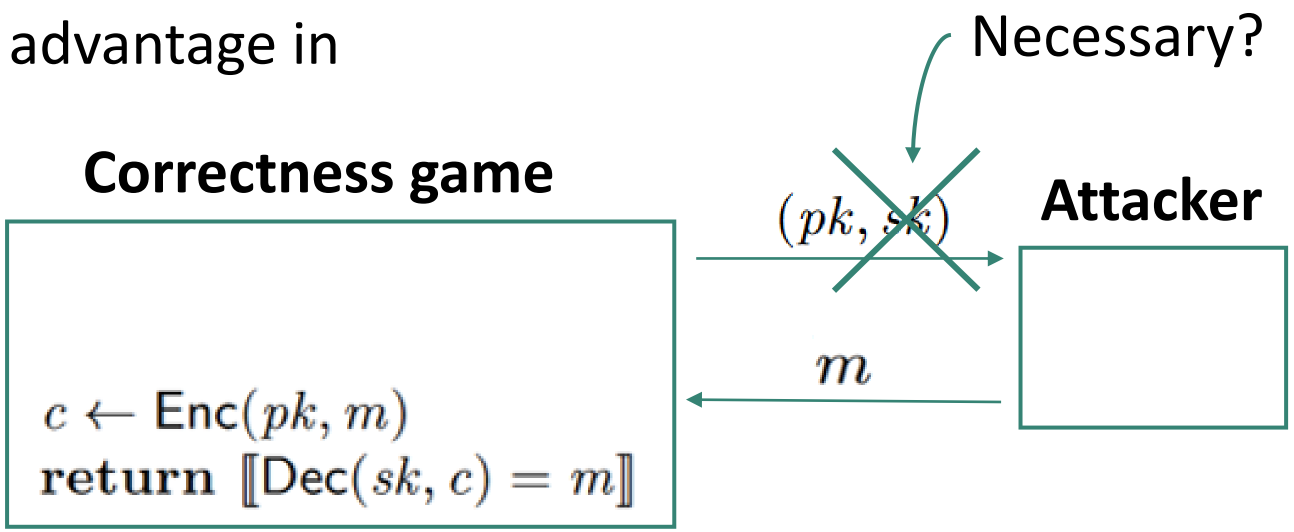

↯ observed by Manuel Barbosa

δ-estimator scripts:

$\triangleq$ advantage in game **without sk**

Applicability issue

Concrete δ – estimations ↯ security proofs

Our results (nutshell)

Tighter bound for FO with explicit rejection ($FO^{\perp}$) for randomised schemes:

→ **Aligns** QROM results for **the two rejection types**

Bounds work with **sk-less failure notions** → **estimator-script-compatible** ☺

Our results

Tighter bound for FO with explicit rejection ($FO^{\perp}$) for randomised schemes:

$$\mathrm{INDCCA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) \;\leq\; \mathrm{INDCPA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) + T_{\mathrm{SPREAD}} + T_{\mathrm{FAIL}}$$

Our results

Tighter bound for FO with explicit rejection ($FO^{\perp}$) for randomised schemes:

$$\mathrm{INDCCA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) \leq \mathrm{INDCPA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) + T_{\mathrm{SPREAD}} + T_{\mathrm{FAIL}}$$

Essentially $4 \cdot \sqrt{\#\ \text{queries} \cdot \mathrm{INDCPA}(\mathrm{PKE})}$

How? Semi-classical One-Way to Hiding (tailored)

Why not double-sided? Same bound

Why not MRM? $4 \cdot \#\ \text{queries}^2 \cdot \mathrm{INDCPA}(\mathrm{PKE})$

Our results

Tighter bound for FO with explicit rejection ($FO^{\perp}$) for randomised schemes:

$$\mathrm{INDCCA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) \leq \mathrm{INDCPA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) + T_{\mathrm{SPREAD}} + T_{\mathrm{FAIL}}$$

$$T_{\mathrm{SPREAD}} = \frac{2^{65} \cdot q}{\sqrt{2^{\gamma}}}$$

γ: PKE spreadness ('entropy')

DFMS22: $\frac{24 \cdot q \cdot \sqrt{q \cdot q_{\mathrm{Decaps}}}}{\sqrt[4]{2^{\gamma}}}$

q: # RO queries

q_{Decaps}: # CCA queries (NIST: 2^{64})

Our results

Tighter bound for FO with explicit rejection ($\mathrm{FO}^{\perp}$) for randomised schemes:

$$\mathrm{INDCCA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) \;\leq\; \mathrm{INDCPA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) + T_{\mathrm{SPREAD}} + T_{\mathrm{FAIL}}$$

Bound **also works for implicit rejection** (due to BH+19).

Conjecture

Implicit: smaller T_{SPREAD} possible

Our results

Tighter bound for FO with explicit rejection ($\mathrm{FO}^{\perp}$) for randomised schemes:

$$\mathrm{INDCCA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) \leq \mathrm{INDCPA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) + T_{\mathrm{SPREAD}} + T_{\mathrm{FAIL}}$$

T_{FAIL}: failure-finding game advantage **without sk**

Previous work:

Implicit: Essentially $8q^2 \cdot \delta$

Explicit: $24 \cdot q^2 \cdot \delta$

$\mathrm{PKE}^{\mathrm{derand}}$: c = Encrypt(pk, m; r), r = $\mathrm{Hash}_{\mathrm{rand}}$(m)

Our results

Tighter bound for FO with explicit rejection ($\mathrm{FO}^{\perp}$) for randomised schemes:

$$\mathrm{INDCCA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) \leq \mathrm{INDCPA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) + T_{\mathrm{SPREAD}} + T_{\mathrm{FAIL}}$$

T_{FAIL}: failure-finding game advantage **without sk**

3 ways to bound:

Failure attacker with CCA oracle, somewhat contrived ROM:

$$T_{\mathrm{FAIL}} = \mathrm{FAILURE} - \mathrm{CCA}\ (\mathrm{PKE}^{\mathrm{derand}})$$

Previous work:

Implicit: Essentially $8q^2 \cdot \delta$

Explicit: $24 \cdot q^2 \cdot \delta$

$\text{PKE}^{\text{derand}}$: c = Encrypt(pk, m; r), r = $\text{Hash}_{\text{rand}}$(m)

Our results

Tighter bound for FO with explicit rejection ($\text{FO}^{\perp}$) for randomised schemes:

$$\text{INDCCA}(\text{FO}^{\perp}(\text{PKE})) \leq \text{INDCPA}(\text{FO}^{\perp}(\text{PKE})) + T_{\text{SPREAD}} + T_{\text{FAIL}}$$

T_{FAIL}: failure-finding game advantage **without sk**

3 ways to bound:

Failure attacker **w'out** CCA oracle, somewhat contrived ROM:

$$T_{\text{FAIL}} = \boldsymbol{q}_{\textbf{Decaps}} \cdot \text{FAILURE} - \text{CPA}\ (\text{PKE}^{\text{derand}})$$

Previous work:

Implicit: Essentially $8q^2 \cdot \delta$

Explicit: $24 \cdot q^2 \cdot \delta$

$\mathrm{PKE}^{\mathrm{derand}}$: c = Encrypt(pk, m; r), r = $\mathrm{Hash}_{\mathrm{rand}}$(m)

Our results

Tighter bound for FO with explicit rejection ($\mathrm{FO}^{\perp}$) for randomised schemes:

$$\mathrm{INDCCA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) \leq \mathrm{INDCPA}(\mathrm{FO}^{\perp}(\mathrm{PKE})) + T_{\mathrm{SPREAD}} + T_{\mathrm{FAIL}}$$

T_{FAIL}: failure-finding game advantage **without sk**

3 ways to bound:

Breaking down $\mathrm{FAILURE} - \mathrm{CPA}\ (\mathrm{PKE}^{\mathrm{derand}})$, **generically**

- in terms of **PKE**, no contrived ROM
- fine-grained term compatible with existing δ-estimator scripts

Previous work:

Implicit: Essentially $8q^2 \cdot \delta$

Explicit: $24 \cdot q^2 \cdot \delta$

$\mathrm{PKE}^{\mathrm{derand}}$: c = Encrypt(pk, m; r), r = $\mathrm{Hash}_{\mathrm{rand}}$(m)

Our results

$\mathrm{FAILURE}-\mathrm{CPA}\left(\mathrm{PKE}^{\mathrm{derand}}\right)=$ sum of two bounds:

- Finding non-generic (key-dependent) failures for PKE
- Finding generic (key-independent) failures for $\mathrm{PKE}^{\mathrm{derand}}$

$\text{PKE}^{\text{derand}}$: c = Encrypt(pk, m; r), r = $\text{Hash}_{\text{rand}}$(m)

Our results

$\text{FAILURE} - \text{CPA}\left(\text{PKE}^{\text{derand}}\right) =$ sum of two bounds:

- **Finding non-generic (key-dependent) failures for** PKE
- Finding generic (key-independent) failures for $\text{PKE}^{\text{derand}}$

NonGenFail game

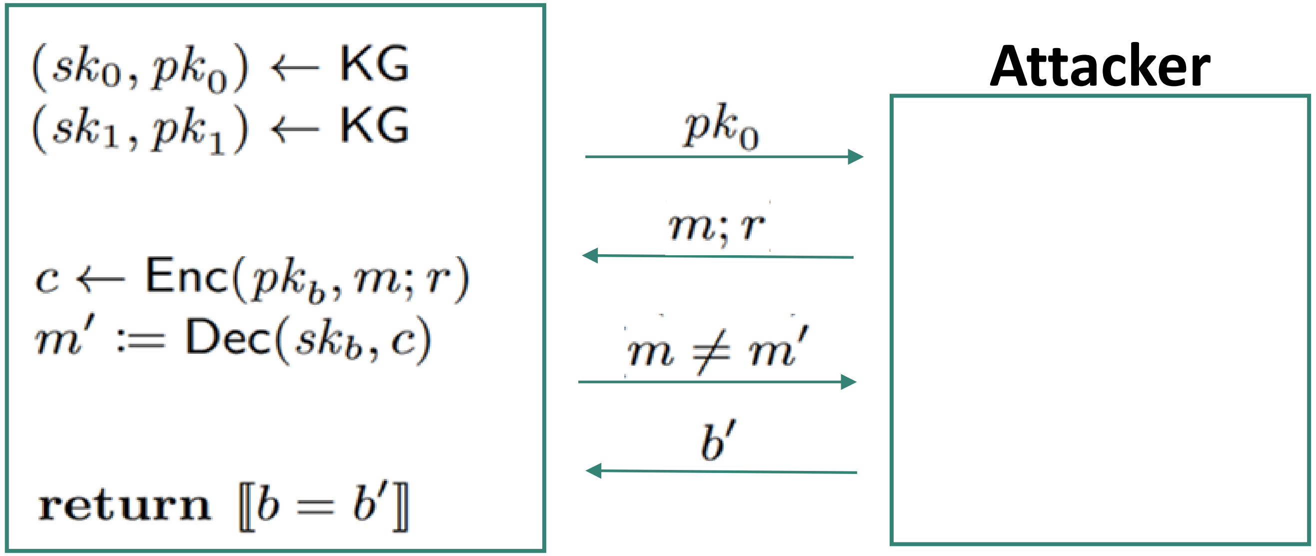

Task: Tell key pairs apart
with single 'does this fail' query

$\mathrm{PKE}^{\mathrm{deran}}$

Conjecture

Lattice-based: NonGenFail ≈ IND-CPA

Our results

$\mathrm{FAILURE}-\mathrm{CPA}\left(\mathrm{PKE}^{\mathrm{derand}}\right)=$ sum of two bounds:

- **Finding non-generic (key-dependent) failures for** PKE
- Finding generic (key-independent) failures for $\mathrm{PKE}^{\mathrm{derand}}$

NonGenFail game

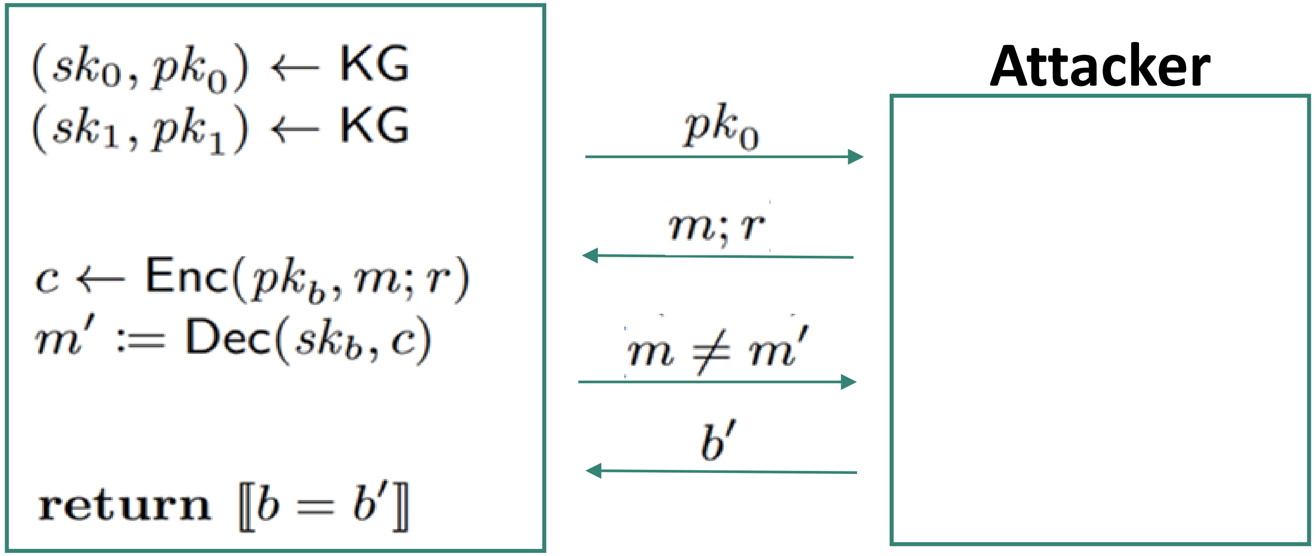

Task: Tell key pairs apart
with single 'does this fail' query

$\mathrm{PKE}^{\mathrm{derand}}$: c = Encrypt(pk, m; r), r = $\mathrm{Hash}_{\mathrm{rand}}$(m)

Our results

$\mathrm{FAILURE} - \mathrm{CPA}\left(\mathrm{PKE}^{\mathrm{derand}}\right) =$ sum of two bounds:

- Finding non-generic (key-dependent) failures
- **Finding generic (key-independent) failures** for $\mathrm{PKE}^{\mathrm{derand}}$

GenFail game

$(pk, sk) \leftarrow \mathsf{KG}$
$c := \mathsf{Enc}(pk, m; \mathsf{Hash}_{\mathsf{rand}}(m))$
$m' := \mathsf{Dec}(sk, c)$
$\mathbf{return}\ [\![m' \neq m]\!]$

m ← **Attacker**

Task: Find m failing for $\mathrm{PKE}^{\mathrm{derand}}$
without even knowing pk

$\mathrm{PKE}^{\mathrm{derand}}$: c = Encrypt(pk, m; r), r = $\mathrm{Hash}_{\mathrm{rand}}$(m)

Our results

$\mathrm{FAILURE-CPA}\left(\mathrm{PKE}^{\mathrm{derand}}\right)$ = sum of two bounds:

- Finding non-generic (key-dependent) failures
- **Finding generic (key-independent) failures** for $\mathrm{PKE}^{\mathrm{derand}}$

GenFail game

Analysis via new QROM ‚find large values‘ bounds

$(pk, sk) \leftarrow \mathsf{KG}$
$c := \mathsf{Enc}(pk, m; \mathsf{Hash}_{\mathsf{rand}}(m))$
$m' := \mathsf{Dec}(sk, c)$
$\textbf{return } [\![m' \neq m]\!]$

m ← **Attacker**

Task: Find m failing for $\mathrm{PKE}^{\mathrm{derand}}$
without even knowing pk

Finding generic failures

'Generic Failure' term = $\tilde{\delta} + T_{\tilde{\delta}}$:

$\tilde{\delta}$:= computed δ- estimate

$T_{\tilde{\delta}} \approx \left(\sqrt{-ln\,(\tilde{\delta})} + \sqrt{ln\,(q_{RO})}\right) \cdot \tilde{\delta}$ if failure tail envelope has Gaussian tail bound

Otherwise:

$T_{\tilde{\delta}} \approx q_{RO}$ · decryption failure rate variance

pessimistic: $< \tilde{\delta}$

Conjecture

Lattice-based: variance very small

Proof technique: Extractable QROM (DFMS22)

Idea: ROM-like reduction via preimage extraction

FO proof:

$O = \mathrm{Hash}_{\mathrm{rand}}: M \to R$

CCA simulation:
Book-keep $\mathrm{Hash}_{\mathrm{rand}}$ queries

Proof technique: Extractable QROM (DFMS22)

Idea: ROM-like reduction via preimage extraction

QROM $O: X \to Y$ via compressed oracle (Zha19)

FO proof:

$O = \mathrm{Hash}_{\mathrm{rand}}: M \to R$

Proof technique: Extractable QROM (DFMS22)

Idea: ROM-like reduction via preimage extraction

QROM $O: X \to Y$ via compressed oracle (Zha19)

+ interface Extract_f for $f: X \times Y \to T$:

$\text{Extract}_f(\text{t})$:

Collapse oracle database such that

- for one x, $f(x, y) = t$ for all y in x's database superposition

Return x

FO proof:

$O = \text{Hash}_{\text{rand}}: M \to R$

Proof technique: Extractable QROM (DFMS22)

Idea: ROM-like reduction via preimage extraction

QROM $O: X \to Y$ via compressed oracle (Zha19)

+ interface $\mathrm{Extract}_f$ for $f: X \times Y \to T$:

$\mathrm{Extract}_f(\mathrm{t})$:

Collapse oracle database such that

- for one x, $f(x, y) = t$ for all y in x's database superposition

Return x

FO proof:

$O = \mathrm{Hash}_{\mathrm{rand}}: M \to R$

$f = \mathrm{Encrypt}: M \times R \to C$

$\mathrm{Extract}_f(\mathrm{c})$ = 'preimage' *m*

Proof technique: Extractable QROM (DFMS22)

Idea: ROM-like reduction via preimage extraction

QROM $O: X \to Y$ via compressed oracle (Zha19)

+ interface $\mathrm{Extract}_f$ for $f: X \times Y \to T$:

$\mathrm{Extract}_f(\mathrm{t})$:

Collapse oracle database such that

- for one x, $f(x, y) = t$ for all y in x's database superposition

Return x

FO proof:

$O = \mathrm{Hash}_{\mathrm{rand}}: M \to R$

$f = \mathrm{Encrypt}: M \times R \to C$

$\mathrm{Extract}_f(\mathrm{c})$ = 'preimage' *m*

'Surprising' ≙ PKE spreadness

$\mathrm{Extract}_f$ commutes nicely with O-operations for sufficiently surprising f.

Proof technique: Extractable QROM (DFMS22)

Idea: ROM-like reduction via preimage extraction

QROM $O: X \to Y$ via compressed oracle (Zha19)

+ interface $\mathrm{Extract}_f$ for $f: X \times Y \to T$:

$\mathrm{Extract}_f(\mathrm{t})$:

Collapse oracle database such that

- for one x, $f(x, y) = t$ for all y in x's database superposition

Return x

FO proof:

$O = \mathrm{Hash}_{\mathrm{rand}}: M \to R$

$f = \mathrm{Encrypt}: M \times R \to C$

$\mathrm{Extract}_f(\mathrm{c})$ = 'preimage' *m*

'Surprising' ≙ PKE spreadness

Contribution: extractable QROM OWTH

$\mathrm{Extract}_f$ commutes nicely with O-operations for sufficiently surprising f.

Conclusion

Thanks for listening!

Eprint: 2022/365

Tighter bound for $\mathrm{FO}^{\perp}$, alternative bound for implicit

$\mathrm{FAILURE} - \mathrm{CCA}\ (\mathrm{PKE}^{\mathrm{derand}})$

↳ $\mathrm{FAILURE} - \mathrm{CPA}\ (\mathrm{PKE}^{\mathrm{derand}})$

↳ $\mathrm{NONGENFAIL}\ (\mathrm{PKE}) + \mathrm{GENFAIL}\ (\mathrm{PKE}^{\mathrm{derand}})$

↳ $\mathrm{sublogarithmic} \cdot \delta-$ estimate

Trade-off: Possibly better bounds vs. less work for concrete schemes

script-compatible

QROM tools: 'large value search' results + proof strategy:

Reduction needs to
- Book-keep queries
- Simulate hash values via $

→ extractable OWTH

Bonus: IND-CCA of FO in the ROM

FO encapsulation

Key: k = Hash_{key}(m), m = $

Ciphertext: r = Hash_{rand}(m)
c = Encrypt(pk, m; r)

FO decapsulation

m' = Decrypt(sk, c)
k = Hash_{key}(m')

IND: Breaking IND = breaking PKE

CCA: Book-keep queries to Hash_{rand}
Look up m encrypting to c
Return Hash_{key}(m)

,PKE entropy'

Decaps simulation fails if:

- c valid, but m not yet queried → γ-spreadness
- c stems from 'failing' m:
  - c = Encrypt(m) with r = Hash_{rand}(m)
  - Decrypt(c) $\neq$ m

Correctness game against derandomised PKE

Advantage < $q_{RO} \cdot \delta$

In the quantum ROM?

FO encapsulation

Key: $k = \text{Hash}_{\text{key}}(m)$, $m = \$$

Ciphertext: $r = \text{Hash}_{\text{rand}}(m)$
$c = \text{Encrypt}(pk, m; r)$

FO decapsulation

$m' = \text{Decrypt}(sk, c)$
$k = \text{Hash}_{\text{key}}(m')$

IND: Breaking IND = breaking PKE

CCA: Simulation that still fails for failing m

Replace $\text{Hash}_{\text{rand}}$ **with 'perfectly correct' oracle**

Advantage < $q_{RO}^2 \cdot \delta$

In the quantum ROM?

☑ One-way to hiding (OWTH)
U14, AHU19, BH+19, KS+21

IND: Breaking IND = breaking PKE

CCA: Book-keep queries to $\text{Hash}_{\text{rand}}$
Look up m encrypting to c
Return $\text{Hash}_{\text{key}}(m)$

FO encapsulation

Key: $k = \text{Hash}_{\text{key}}(m)$, $m = \$$

Ciphertext: $r = \text{Hash}_{\text{rand}}(m)$
$c = \text{Encrypt}(pk, m; r)$

FO decapsulation

$m' = \text{Decrypt}(sk, c)$
$k = \text{Hash}_{\text{key}}(m')$

This work

ROM-like simulation via extractable QROM
+
OWTH in extractable QROM

Bonus: tail bound of failure tail envelope

$T_{\tilde{\delta}} \approx \left(\sqrt{-ln\left(\tilde{\delta}\right)} + \sqrt{ln\left(q_{RO}\right)}\right) \cdot \tilde{\delta}$ if failure tail envelope has Gaussian tail bound

Failure tail envelope: $\tau(\mathrm{t}) \coloneqq \max_{m} \Pr_{r}\left[\Pr_{pk,sk}[m, r \text{ fail for } pk, sk] \geq t\right]$

Gaussian tail bound: $\tau(\mathrm{t}) \leq \exp\left(-\frac{1}{\tilde{\delta}^2} \cdot (t - \delta_{ik})^2\right)$

δ_{ik}: $\max_{m} \Pr_{r,pk,sk}[m, r \text{ fail for } pk, sk]$

Bonus: Compressed oracle (Zha19)

- Oracle database initalised to $D := \bigotimes_{x \in query\ domain} |x, \perp >_{D_x}$
- Process queries $|x, y>$ by applying
  - F_{D_x} to output register of D_x

$$F_{D_x}|\psi > := \begin{cases} uniform\ superposition, & |\psi > = \perp \\ \perp, & |\psi > = uniform\ superposition \\ |\psi >, & |\psi > orthogonal\ to\ \perp, uniform \end{cases}$$

  - $\mathrm{CNOT}^{\otimes}_{D_x:Y}$ to D_x, query output register Y
  - F_{D_x} to output register of D_x