

Adding functionality to post-quantum cryptography with variants of the Fujisaki–Okamoto transform

Douglas Stebila
University of Waterloo

Based on a series of joint work with [Lewis Glabush](#) (C&O MMath → EPFL PhD), [Felix Günther](#) (IBM Research Zürich), [Britta Hale](#) (Naval Postgraduate School), [Kathrin Hövelmanns](#) (TU Eindhoven), and [Elizabeth Van Oorschot](#) (C&O URA, McGill).

We acknowledge the support of NSERC.

Some motivating attacks

Public-key encryption

Alice generates a key pair and publishes her public key:


$$\text{KeyGen} \rightarrow (pk, sk)$$

Anyone can encrypt a message to Alice:

$$\text{Enc}(pk, m) \rightarrow c$$

Only Alice can decrypt a ciphertext:

$$\text{Dec}(sk, c) \rightarrow m$$

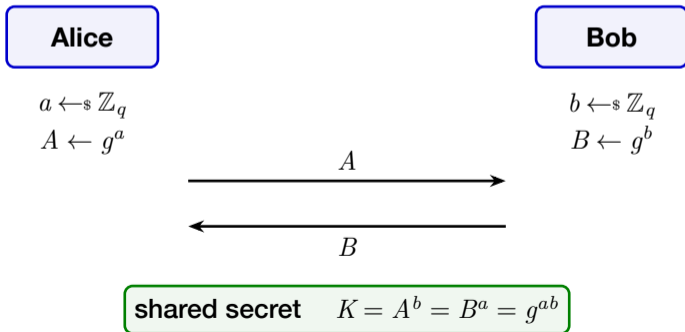


...but public key encryption is a polite fiction

We almost never encrypt application data with public-key encryption.

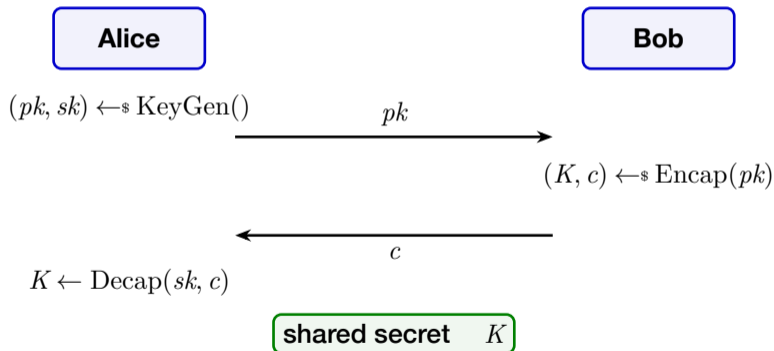
Instead: use asymmetric algorithms to agree on a **shared secret**, then encrypt with fast symmetric crypto.

e.g. Diffie-Hellman key exchange (in a group $\langle g \rangle$ of order q):



Key encapsulation mechanisms (KEMs)

A generalization of Diffie–Hellman key exchange:



KEMs are at the heart of the post-quantum transition

Post-quantum KEM standardization:

- ML-KEM* – NIST standard, 2024, based on module lattices
- HQC – NIST selected, 2025, based on error-correcting codes
- FrodoKEM* – ISO standard, 2026, based on lattices

Adoption:

- TLS (the web)
- Signal, SSH, iMessage, ...


How to build a KEM?

Every standardized post-quantum KEM is built via the **Fujisaki-Okamoto (FO) transform**.



What does it mean for a KEM to be secure?

A KEM is secure if its key is **indistinguishable from random**, even to an adversary who sees the public key and ciphertext.




A KEM is **IND-CPA-secure** if the adversary's advantage is small:

$$\text{Adv}_{\text{KEM}}^{\text{IND-CPA}}(\mathcal{A}) = \left| \text{Prob}[b' = b] - \frac{1}{2} \right|$$

Active attackers: IND-CCA (chosen ciphertext attacks)

Real attackers don't just watch — they **inject and modify** ciphertexts and observe what the receiver does. So we give the adversary a **decapsulation oracle**:



Indistinguishability against **chosen ciphertext attacks**

Note

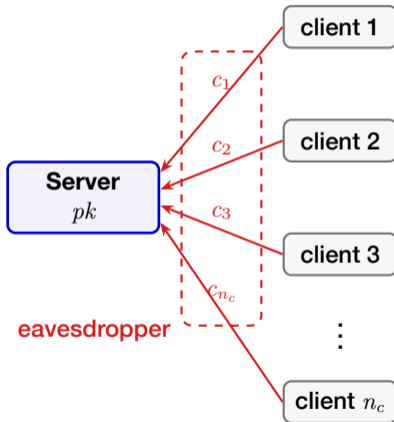
For IND-CCA, adversary is given **one challenge ciphertext** to break.

Attack 1 – the setting

One server, one key pair, *many* clients.

The adversary passively collects n_c encapsulation ciphertexts off the wire.

Goal: recover the shared key of *any one* of them.




FO makes encryption a function of the message

Inside the FO transform, encryption is **derandomized**:

$$c = \text{Enc}(pk, m; G(m)) \quad (\text{the randomness of Enc is } G(m), \text{ a hash of } m).$$

So c is a deterministic, *public* function of m — anyone can recompute it.

⇒ If the message space is small, an eavesdropper can just **guess and check**.



- challenge ciphertexts
- adversary's N guesses
- collision = a recovered key

A pure birthday / collision argument:

$$\text{Prob}[\text{collision}] \approx \frac{N n_c}{|\mathcal{M}|}.$$

For FrodoKEM-640 & HQC-128:

$$|\mathcal{M}| = 2^{128}$$

Take $N = n_c = 2^{64}$:

high success probability.

Multi-ciphertext security is halved

128-bit single-challenge security
→ ~64-bit **multi-target** security.

Identified by NIST & Bernstein against
FrodoKEM-640 and HQC-128.


FO decapsulation re-encrypts to check the ciphertext

Attack 2:
implementation bug

After derandomization, each message has *exactly one* valid ciphertext:

$$c = \text{Enc}(pk, m; G(m)).$$

So on receipt, the FO transform doesn't just decrypt; it **re-encrypts and checks**.



Idea: if the ciphertext was well-formed, the adversary “knew” the message, so decrypting it doesn't help the adversary. This check is what upgrades *passive* IND-CPA security to *active* IND-CCA security: a malformed ciphertext fails the check and is rejected, so the **decapsulation oracle** leaks nothing.

Attack 2: a missing re-encryption check

Attack 2:
implementation bug

Remove that one check...

...and decapsulation will **accept malformed ciphertexts** — the KEM silently falls back to mere passive (CPA) security.

An **active attacker** feeds clever ciphertexts to the decapsulation oracle and tries to learn the secret: exactly the **chosen-ciphertext attack** the check was there to stop.

This is not hypothetical:

For **19 months**, HQC's reference implementation effectively skipped the re-encryption check — and basic correctness tests still passed.

- Two interacting bugs in HQC's clever hand-crafted constant-time code bypassed the check.
- *Honest* ciphertexts still decapsulated correctly – so functional tests passed.
- The broken code was picked up downstream (e.g. liboqs).

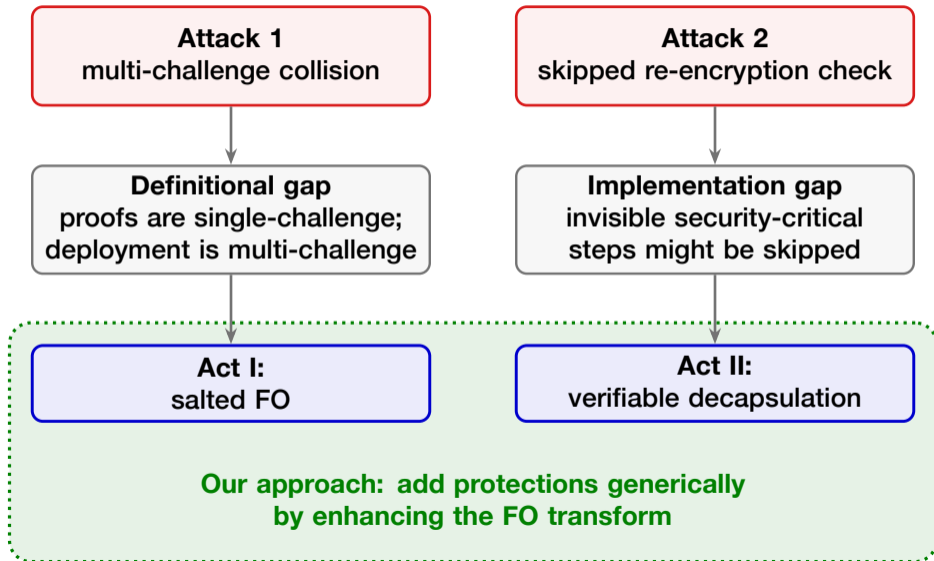
We've seen this pattern before:

Echoes Apple's "goto fail" (2014): one skipped verification step in TLS that was invisible to ordinary testing.

A lesson:

Danger from security-critical steps with no functional footprint.

Observations



Outline

The FO transform

Act I: Multi-target security

Act II: Verifiable decapsulation

Wrapping up

The Fujisaki–Okamoto transform in detail


What FO promises



- Start from a PKE with only a **weak** guarantee (no protection against active attackers).
- Obtain a KEM secure against **chosen-ciphertext** attacks.

The modular view: decompose $FO = U \circ T$

Due to Hofheinz–Hövelmanns–Kiltz (2017): analyze FO as two composable steps, which perform two different roles.



The T-transform: derandomization + re-encryption check

1. Derandomize:

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

where G is a random hash function.

Each message now has *one* canonical ciphertext.

2. Re-encryption check:

1. **Decrypt:** $m' \leftarrow \text{Dec}(sk, c)$
2. **Re-encrypt:** $c' \leftarrow \text{Enc}_1(pk, m')$
3. **Check:** $c = c'$

The scheme becomes **rigid**: a ciphertext is accepted only if it is the genuine encryption of what it decrypts to.

The T-transform: rigidity

Rigidity

A ciphertext is accepted only if it is the genuine encryption of what it decrypts to.

Why does rigidity matter?

- It makes the decapsulation oracle *useless* to an attacker — it only ever decrypts honestly-formed ciphertexts.
- This is what turns *passive* (CPA) into *active* (CCA) security.

Closely connected to our motivating attacks

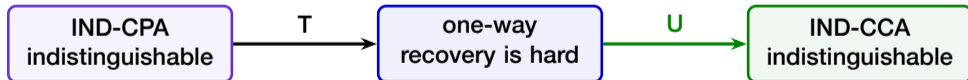
- Skipping rigidity check enabled the HQC attack.
- Determinism enabled FrodoKEM multi-ciphertext attack.

One-wayness vs. indistinguishability

Determinism costs us indistinguishability

A *deterministic* encryption can never satisfy indistinguishability (IND-CPA): to test whether c encrypts a candidate m , just recompute $\text{Enc}_1(pk, m)$ and compare.

So the strongest goal we can ask of T 's output is **one-wayness**: given a ciphertext for a random message, it is hard to recover the message.



The U transform will lift us back up to indistinguishability.

The U-transform: one-way PKE \rightarrow indistinguishable KEM

Encapsulation:

1. Pick random m
2. Encrypt: $c \leftarrow \text{Enc}_1(pk, m)$
3. Compute shared secret:
 $K \leftarrow H(m, c)$

Decapsulation of c :

1. Decrypt: $m' \leftarrow \text{Dec}_1(sk, c)$
(with re-encryption check)
2. Compute shared secret:
 $K \leftarrow H(m', c)$

Why does hashing give us indistinguishability?

- Because the message is hard to *recover* (one-wayness), the attacker can never *query* the random oracle on it.
- So output $H(m, c)$ is a fresh, uniformly random value to the adversary.

The U-transform: one-way PKE \rightarrow indistinguishable KEM

What to do with an invalid ciphertext during decapsulation?

U^\perp : explicit rejection


1. **Decrypt:** $m' \leftarrow \text{Dec}_1(sk, c)$
(with re-encryption check)
2. **If $m' = \perp$, return \perp**
3. **Else: Compute shared secret:**
 $K \leftarrow H(m', c)$

U^\neq : implicit rejection

1. **Decrypt:** $m' \leftarrow \text{Dec}_1(sk, c)$
(with re-encryption check)
2. **If $m' = \perp$, return $H'(z, c)$ for hidden secret z**
3. **Else: Compute shared secret:**
 $K \leftarrow H(m', c)$

Every standardized PQ KEM
uses **implicit rejection**.

Summary of the FO transform



Theorem (FO security informally)

If the PKE is IND-CPA-secure, then the KEM $FO = U \circ T$ is IND-CCA-secure, in the (quantum) random oracle model.

- Standardized PQ KEMs ML-KEM, FrodoKEM, HQC each use a different FO variant, but all follow the overall T-then-U pattern.


Attack #1: FO yields single-challenge security, not multi-challenge security.

Attack #2: FO only provides security if re-encryption check done properly.

Act I: Multi-target security and the salted FO transform

Lewis Glabush, Kathrin Hövelmanns, Douglas Stebila. On the multi-target security of post-quantum KEMs. *IACR Communications in Cryptology*, 3(1), May 2026.

The definitional gap



One can give a generic hybrid argument from single-challenge to multi-challenge, but that loses a factor of $n_u n_c$ which appreciably decreases security.

Can we do better?

The multi-ciphertext attack, reviewed

Adversary is given:

Public key pk

n_c challenge ciphertexts c_1, \dots, c_{n_c}
encrypting random messages m_i :

$$c_i \leftarrow \text{Enc}_1(pk, m_i) = \text{Enc}(pk, m_i; G(m_i))$$

The attack:

1. Sample N messages m'_1, \dots, m'_N ;
2. Encrypt each (*this is possible!*
after T, encryption is
deterministic and pk is public)

$$\begin{aligned}c'_j &\leftarrow \text{Enc}_1(pk, m'_j) \\ &= \text{Enc}(pk, m'_j; G(m'_j))\end{aligned}$$

3. Store (c'_j, m'_j)
4. Look for a match with a
challenge:
if $c'_j = c_i$, then $m_i = m'_j$

The multi-ciphertext attack, reviewed

$$\text{Prob}[\text{successful attack}] \approx \frac{N n_c}{|\mathcal{M}|}.$$

For $|\mathcal{M}| = 2^{128}$, given $n_c = 2^{64}$ challenge ciphertexts, an adversary who does $N = 2^{64}$ work succeeds with high probability.

(This doesn't affect single challenge security $n_c = 1$.)


How to fix it?

Idea 1: increase $|\mathcal{M}|$

- 😞 For FrodoKEM, increasing message size leads to much bigger parameters

Idea 2: increase unpredictability with a random public salt


The fix: add a public salt in the T transform



Salted FO (SFO):

1. Pick random m
2. **Sample a fresh random salt s**
3. **Encrypt:** $c \leftarrow \text{Enc}(pk, m; G(\mathbf{s}, m))$
4. **Transmit (c, \mathbf{s})**
5. **Compute shared secret:**
 $K = H(m, c, \mathbf{s})$.

The fix: add a public salt in the T transform




For a collision to help, the adversary must now also **match the salt** — which is information-theoretically random. Hard to guess in advance; defeats batch and precomputation attacks.

$$\text{Prob}[\text{successful attack}] \approx \frac{N n_c}{|\mathcal{M}| 2^{\text{len}_s}}.$$

Bottom line: if $\text{len}_s \geq 256$ (32 bytes of extra communication), we thwart this multi-ciphertext attack.

Wait, there's more: an almost-free collision



- **challenge ciphertexts** c_1, \dots, c_{n_c}
for messages selected randomly
(with replacement)

$$\text{Prob}[\text{collision}] \approx \frac{n_c^2}{|\mathcal{M}|}$$

If two challenges encrypt the same message, they lead to the same key and ciphertext:

$$m_i = m_j \Rightarrow c_i = c_j, K_i = K_j$$

But *independent* random keys would not match.

\Rightarrow Adversary can distinguish real vs. random challenge set **without any queries!**

Salts fix this too!

Salts reduce these types of collisions:

$$\text{Prob}[\text{collision}] \approx \frac{n_c^2}{|\mathcal{M}| 2^{\text{len}_s}}$$

Salted FO results

Theorem: Salted FO is tightly multi-target secure in the classical random oracle model.

$$\text{Adv}_{n_c, n_u}^{\text{IND-CCA}}(\mathcal{A}) \leq \text{Adv}_{n_c, n_u}^{\text{IND-CPA}}(\mathcal{B}) + \frac{n_u n_c^2}{|\mathcal{M}| 2^{\text{len}_s}} + \frac{9 q_{\text{RO}}}{|\mathcal{M}|} + q_{\text{RO}} \cdot \delta.$$

multi-ciphertext (n_c), multi-user* (n_u), δ -multi-user correctness

Tight reduction: multi-target IND-CPA of PKE \implies multi-target IND-CCA of KEM.

The dangerous $n_u n_c$ -type terms are now divided by 2^{len_s} — the salt buys back security.

What about quantum attackers?

A quantum random oracle model (QROM) proof exists — but is *non-tight*.

*Multi-user security addressed by hashing the public key into shared secret derivation: Duman, Hövelmanns, Kiltz, Lyubashevsky, Seiler. Faster lattice-based KEMs via a generic Fujisaki–Okamoto transform using prefix hashing. *ACM CCS* 2021.

Multi-target security, salted vs. unsalted

Unsalted FO	$\frac{n_u n_c^2 + (2n_c + 1) q_{\text{RO}}}{ \mathcal{M} }$	≈ 1 in the FrodoKEM / HQC regime
Salted FO	$\frac{n_u n_c^2}{ \mathcal{M} 2^{\text{len}_s}} + \frac{9 q_{\text{RO}}}{ \mathcal{M} }$	negligible for salt length ≈ 32 bytes

Bottom line

At the 128-bit level, replacing FO with Salted FO recovers ~ 62 bits of multi-target security for FrodoKEM and HQC.

Cost: one extra public salt per ciphertext; $< 1\%$ runtime / bandwidth overhead.

FrodoKEM ISO standard uses Salted FO.

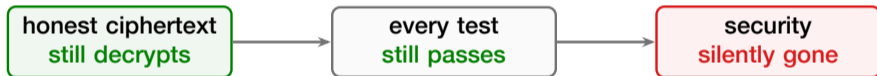
Act II: Verifiable decapsulation

Lewis Glabush, Felix Günther, Kathrin Hövelmanns, Douglas Stebila. Verifiable decapsulation: recognizing faulty implementations of post-quantum KEMs. *CRYPTO 2025*.

Lewis Glabush, Felix Günther, Britta Hale, Elizabeth Van Oorschot, Douglas Stebila. Augmenting post-quantum KEMs for verifiable decapsulation with rejection hiding. Under review, 2026.

Brittle cryptography

- HQC's re-encryption check was effectively skipped – unnoticed for **19 months**.
- Why was it missed?
 - Skipping it breaks nothing visible: honest ciphertexts decrypt, all tests pass.
 - Only the security is lost.



Principle (Heninger 2019; Fischlin–Günther 2023)

Tie security to functionality: a faulty-but-benign implementation should break something *visible*.

“Verifiable verification” and confirmation codes

Fischlin and Günther’s **Verifiable verification**:

Modify the scheme so its output changes if security-critical verification steps are skipped.

Algorithm produces a **confirmation code** cd :


- An extra deterministic output, **unpredictable** if a security-critical computation is omitted.
- Skip the step
 \implies unlikely to produce a valid cd .
- Formal notion: *confirmation-code unpredictability* (cUP).

Confirmation-code unpredictability advantage doesn’t have to be cryptographically good

Probability $\frac{1}{2}$ suffices: a faulty implementation that fails half the time will still be noticed in basic interop testing.

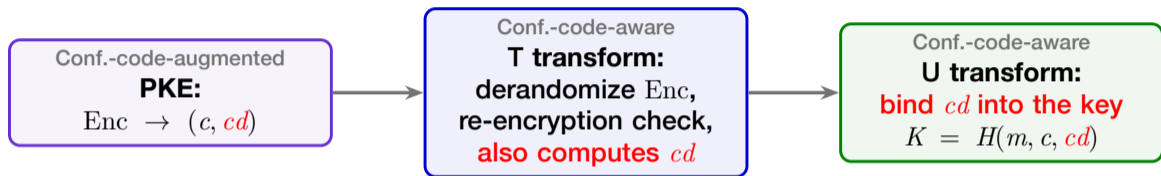
Confirmation code for ML-KEM

- ML-KEM ciphertexts are compressed versions of polynomials in $\mathbb{Z}_{3329}[X]/(X^{256} + 1)$
- **Coefficient-wise compression** rounds \mathbb{Z}_{3329} to \mathbb{Z}_{2^d} ($d = 4, 5, 10, \text{ or } 11$)
- Not exactly bit truncation (since q odd), but intuitively similar:



- Remainder is unpredictable without re-doing most of the encryption computation.
 - Proof technique: model pseudorandom generator in random oracle model.
- Since the remainder is computed during encryption, then thrown away, keeping 12 – 20 bytes for cd is effectively free.

Confirmation code-augmented FO transform



cd is never transmitted: decapsulation can only recompute it by actually re-encrypting.

Still IND-CCA-secure.

Faulty decapsulation implementation (skipping re-encryption and hence code re-computation) now yields a wrong key (**faulty correctness (fCOR)**).

Formalization

Confirmation-code unpredictability:

Without doing the security-critical computation, an adversary cannot produce a valid code cd .

Faulty correctness:

An implementation that skips the security-critical step produces the *wrong* shared secret with noticeable probability.

Theorem: Confirmation-code-augmented FO is secure.

Let $\text{FOC} = \text{UC} \circ \text{TC}$ applied to PKE.

- If PKE is IND-CPA-secure, then FOC is IND-CCA-secure. (ROM, QROM)
- If PKE is cUP-secure, then FOC is fCOR-secure. (ROM)

Application to HQC and ML-KEM

Scheme	Confirmation code <i>cd</i>	Computational overhead	Guessing probability
HQC	1 byte from end of pre-truncation ciphertext	0.25%	0.03
ML-KEM	12 – 20 bytes of discarded compression bits	1.2 – 3.4%	0.17 – 0.68

Catching the original bug

With a confirmation code, HQC's reference-implementation bug surfaces on the **first** interop test, not 19 months later.

We've proposed to NIST that this be included in the final HQC spec.
(Too late for ML-KEM.)

One drawback

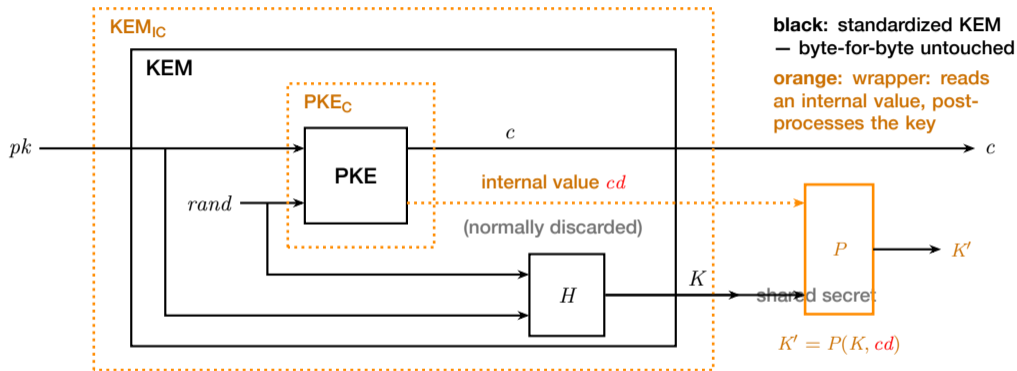
Hashing cd into the key changes the scheme

- Shared secret is different from standard since it now depends on cd
- Hard sell after standardization.

Our solution:

- KEM computations left untouched – as standardized.
- Wrapper reads *internal* values to pull out confirmation code.
- Post-process original shared secret + confirmation code.

Wrapping the standardized KEM



Variants: implicit versus explicit rejection, whether confirmation code is considered internal or public.

This figure shows encapsulation for **internal-code** variant.

Rejection hiding

We had to make sure the confirmation codes didn't undermine the implicit or explicit rejection goals of U^{\neq} and U^{\perp} .

But... what does implicit rejection actually mean?

All major PQC KEMs (ML-KEM, HQC, FrodoKEM, Classic McEliece) use implicit rejection, but it was **never formally defined**.

We define **rejection hiding** and prove that Hofheinz–Hövelmanns–Kiltz's original FO^{\neq} transform satisfies it, plus our confirmation-code-augmented variants.

Wrapping up

Two themes

The multi-target gap

Standard definitions don't necessarily match deployment reality.

Approach: a structural tweak (a **salt**) with a *tight* reduction.

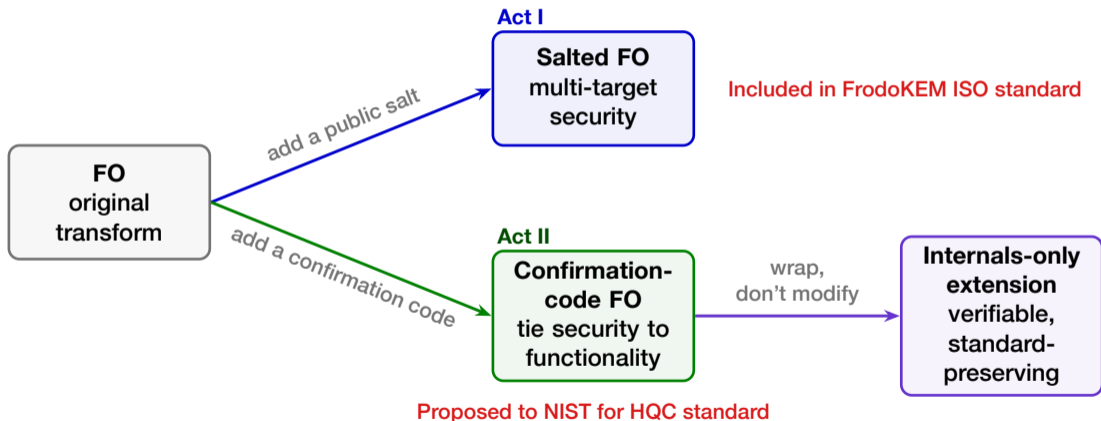
The implementation gap

Implementations can have flaws that may affect security.

Approach: make security-critical steps **functionally observable**.

Both are about closing the distance between what is proven and how it gets used.

Additional functionality via the Fujisaki–Okamoto transform



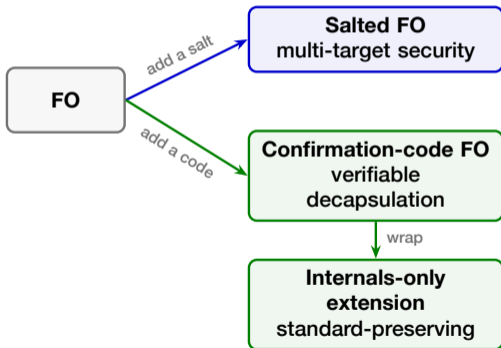
Open questions

- **Tight QROM bounds** for salted FO.
- **Lower bounds** on salt length and confirmation-code length.
- A **combined** multi-target + verifiable variant.
- Verifiable decapsulation for **explicitly-rejecting** deterministic PKEs — no proofs known.
- Re-encryption can create side-channel vulnerabilities – can we get rid of re-encryption altogether? Seems hard.[†]

[†] Hövelmanns, Hülsing, Majenz, Sisinni. (Un)breakable curses - re-encryption in the Fujisaki-Okamoto transform. *Eurocrypt* 2025.

Adding functionality to post-quantum cryptography with variants of the Fujisaki–Okamoto transform

Douglas Stebila, University of Waterloo



- A public **salt** restores multi-target security (now in FrodoKEM ISO standard)
- **Confirmation codes** expose faulty implementations (proposed for HQC)
- Apply generically via variants of FO transform
- Closing the gap between **what is proven** and **how it gets used**

Glabush, Hövelmanns, Stebila. On the multi-target security of post-quantum KEMs. *IACR Communications in Cryptology*, 3(1), May 2026.

Glabush, Günther, Hövelmanns, Stebila. Verifiable decapsulation: recognizing faulty implementations of post-quantum KEMs. *CRYPTO* 2025.

Glabush, Günther, Hale, Van Oorschot, Stebila. Augmenting post-quantum KEMs for verifiable decapsulation with rejection hiding. Under review, 2026.

Appendix

Anatomy of the HQC bug

Bug 1: re-encryption reads the public key from the wrong offset in sk

```
- const uint8_t *pk = sk + SEED_BYTES;  
+ const uint8_t *pk = sk + SEED_BYTES + VEC_K_SIZE_BYTES;
```

Re-encrypts under garbage \Rightarrow the comparison with (u, v) *always fails*.

Bug 2: the constant-time selection mask is inverted

```
- result = (uint8_t) (-((int16_t) result) >> 15);  
+ result -= 1;  
  mc[i] = (m[i] & result) ^ (sigma[i] & ~result);
```

Old mask selects m (= accept) exactly when the comparison *fails*.

Two wrongs make a silent right

Always-failing check \times inverted acceptance = **always accept**. Either bug alone fails correctness tests instantly; together they pass every functional test — only the security is gone.

Implicit vs. explicit rejection — security history

- **Implicit rejection:** proofs available first (HHK 2017), and tighter.
- **Explicit rejection:** early treatments added a key-confirmation hash; first QRROM proofs (Don–Fehr–Majenz–Schaffner 2021) had bigger tightness loss.
- Hövelmanns–Hülsing–Majenz 2022 (“Failing gracefully”): bounds closer to implicit, for *probabilistic* PKEs with certain correctness properties.
- Hövelmanns–Kudinov (eprint 2025/062): explicit rejection essentially as secure as implicit, in the extractable QRROM.
- **Remaining gap:** *deterministic* PKE + explicit rejection — no proofs known (the open question from the closing slide).

Multi-target IND-CCA definition

$\text{Exp}_{\text{KEM}}^{\text{IND}_{n_c, n_u}\text{-CCA}}(\mathcal{A})$	CHALL_j // At most n_c queries
01 $b \leftarrow_{\$} \{0, 1\}$	07 $(c, \text{ss}_0) \leftarrow_{\$} \text{Encaps}(pk_j)$
02 for $j = 1, \dots, n_u$	08 $\text{ss}_1 \leftarrow_{\$} \mathcal{K}$
03 $(pk_j, sk_j) \leftarrow_{\$} \text{Gen}()$	09 $\mathcal{L}_{C_j} \leftarrow \mathcal{L}_{C_j} \cup \{c\}$
04 $\vec{pk} = (pk_1, \dots, pk_{n_u})$	10 return (c, ss_b)
05 $b' \leftarrow \mathcal{A}^{\text{DECAPS}, \text{CHALL}_1, \dots, \text{CHALL}_{n_u}}(\vec{pk})$	$\text{DECAPS}(j, c \notin \mathcal{L}_{C_j})$
06 return $\llbracket b = b' \rrbracket$	11 return $\text{Decaps}(sk_j, c)$

- n_u independent key pairs; n_c challenge ciphertexts per user.
- Adversary gets decapsulation oracles for all users (except on challenge ciphertexts).
- Wins by distinguishing *any* challenge key from random.

The salted FO bound, general form

$$\text{Adv}_{n_c, n_u}^{\text{IND-CCA}}(\mathcal{A}) \leq \text{Adv}_{n_c, n_u}^{\text{IND-CPA}}(\mathcal{B}) + \frac{n_u n_c (n_c - 1)}{|\mathcal{M}| 2^{\text{len}_s}} + \frac{(2t + 1) q_{\text{RO}}}{|\mathcal{M}|} + q_{\text{RO}} \cdot \delta(n_u)$$

- q_{RO} : total number of random-oracle (hash) queries by the adversary.
- $\delta(n_u)$: multi-user correctness error of the PKE.
- t : small constant from the proof; $t \leq 4$ for FrodoKEM-640 parameters, giving the “9” on the main slide.

The QROM bound for salted FO

In the quantum random oracle model, the adversary may query the hash function *in superposition*. Current proof techniques lose square-root factors.

$$\begin{aligned} \text{Adv}_{\text{KEM}}^{\text{IND}_{n_c, n_u}\text{-CPA}}(\mathcal{B}) &\leq 4\sqrt{d \cdot \text{Adv}_{\text{PKE}}^{\text{IND}_{n_c, n_u}\text{-CPA}}(\mathcal{A})} + \frac{n_u n_c (n_c - 1)}{|\mathcal{M}| \cdot 2^{\text{len}_{\text{salt}}}} \\ &\quad + 4\sqrt{\frac{dt(q_H + q_G)}{|\mathcal{M}|}}, \end{aligned}$$

Whether the quadratic loss reflects a real quantum attack or is an artifact of the proof technique is open — for salted FO and for plain FO alike.

cUP and fCOR definitions

Confirmation-code unpredictability (cUP)

$\text{Exp}_{\text{PKE}_C, F, \max}^{\text{cUP}}(\mathcal{A})$

```
01 (pk, sk)  $\leftarrow$  $ KeyGenF()
02 (m, r)  $\leftarrow$  $  $\mathcal{M} \times \mathcal{R}$ 
03 (c, cd)  $\leftarrow$  EncCF(pk, m; r)
04 cd'  $\leftarrow$   $\mathcal{A}^{\bar{F}}$ (pk, sk, c, m, r)
05 return  $\llbracket \text{cd} = \text{cd}' \rrbracket$ 
```

$\bar{F}(x)$ //at most max many queries

```
06 return F(x)
```

Faulty implementation correctness (fCOR)

$\text{Exp}_{\text{KEM}, F, \max}^{\text{fCOR}}(\mathcal{A})$

```
01 (pk, sk)  $\leftarrow$  $ KEM.KeyGenO, F()
02 (c, K0)  $\leftarrow$  $ KEM.EncapsO, F(pk)
03 K'0  $\leftarrow$  KEM.DecapsO, F(sk, c)
04 if K'0  $\neq$  K0 //correctness error
05   return 1 // is a trivial win
06 K1  $\leftarrow$   $\mathcal{A}^{\text{O}, \bar{F}}$ (pk, sk, c)
07 return  $\llbracket K_0 = K_1 \rrbracket$ 
```

$\bar{F}(x)$ //at most max many queries

```
08 return F(x)
```

Confirmation codes for ML-KEM and HQC: the numbers

ML-KEM: confirmation-code guessing probabilities

Scheme	Conditional entropy		Conf. code guessing probability	
	of e_1 coeff. (h_{e_1})	of e_2 coeff. (h_{e_2})	Short ($ S = 2$) (p_2)	Full ($ S = 256$) (p_{256})
<i>Assuming uniform z</i>				
ML-KEM-512	1.2733	2.0188	0.1712	2^{-325}
ML-KEM-768	1.2733	2.0188	0.1712	2^{-325}
ML-KEM-1024	0.6407	2.0069	0.4114	2^{-164}
<i>Worst-case analysis</i>				
ML-KEM-512	0.3219	0.3219	0.6806	2^{-71}
ML-KEM-768	0.3219	0.3219	0.6806	2^{-71}

ML-KEM & HQC: performance overhead

Algorithm		Short conf. code		Full-length conf. code	
		Size (bytes)	Slowdown	Size (bytes)	Slowdown
ML-KEM-512	Encaps	12	3.4%	1152	23.6%
	Decaps		3.1%		23.6%
ML-KEM-768	Encaps	16	2.1%	1536	20.5%
	Decaps		2.1%		15.0%
ML-KEM-1024	Encaps	20	1.2%	1920	14.1%
	Decaps		1.6%		12.2%
HQC-128 with bug	Encaps	1	0.09%	—	—
	Decaps		0.25%		
HQC-128 bug fixed	Encaps	1	0.13%	—	—
	Decaps		0.19%		

Internals-only conf.-code-augmented FO transform variants

explicit rejection

implicit rejection

internal code

Internal code explicit rejection: $UC_{IC}^{\perp} / FOC_{IC}^{\perp}$

Internal code implicit rejection: $UC_{IC}^{\perp} / FOC_{IC}^{\perp}$

KeyGen()	Decaps(sk, c)
01 $(pk, sk) \leftarrow \text{PKE}_C.\text{KeyGen}()$	07 $m' \leftarrow \text{Dec}(sk, c)$
02 return (pk, sk)	08 (c', cd')
<u>Encaps(pk)</u>	$\leftarrow \text{Enc}_C(pk, m'; G(m'))$
03 $m \leftarrow \mathcal{M}$	09 $K' \leftarrow H(m', \text{ID}(pk))$
04 $(c, cd) \leftarrow \text{Enc}_C(pk, m'; G(m))$	10 if $m' = \perp$ or $c' \neq c$ then
05 $K \leftarrow H(m, \text{ID}(pk))$	11 return (\perp, \perp)
06 return (K, c, cd)	12 return (K', cd')

KeyGen()	Decaps($sk = (sk', z), c$)
13 $(pk, sk') \leftarrow \text{PKE}_C.\text{KeyGen}()$	21 $m' \leftarrow \text{Dec}(sk', c)$
14 $z \leftarrow \mathcal{M}$	22 (c', cd')
15 $sk \leftarrow (sk', z)$	$\leftarrow \text{Enc}_C(pk, m'; G(m'))$
16 return (pk, sk)	23 $K' \leftarrow H(m', \text{ID}(pk))$
<u>Encaps(pk)</u>	24 $K'' \leftarrow H(z, c)$
17 $m \leftarrow \mathcal{M}$	25 if $m' = \perp$ or $c' \neq c$ then
18 $(c, cd) \leftarrow \text{Enc}_C(pk, m'; G(m))$	26 return (K'', \perp)
19 $K \leftarrow H(m, \text{ID}(pk))$	27 return (K', cd')
20 return (K, c, cd)	

public code

Public code explicit rejection: $UC_{PC}^{\perp} / FOC_{PC}^{\perp}$

Public code implicit rejection: $UC_{PC}^{\perp} / FOC_{PC}^{\perp}$

KeyGen()	Decaps(sk, c)
01 $(pk, sk) \leftarrow \text{PKE}_C.\text{KeyGen}()$	08 $m' \leftarrow \text{Dec}(sk, c)$
02 return (pk, sk)	09 (c', cd')
<u>Encaps(pk)</u>	$\leftarrow \text{Enc}_C(pk, m'; G(m'))$
03 $m \leftarrow \mathcal{M}$	10 $K' \leftarrow H(m', \text{ID}(pk))$
04 $(c, cd) \leftarrow \text{Enc}_C(pk, m'; G(m))$	11 $cd' \leftarrow J(m', cd')$
05 $K \leftarrow H(m, \text{ID}(pk))$	12 if $m' = \perp$ or $c' \neq c$ then
06 $cd \leftarrow J(m, cd)$	13 return (\perp, \perp)
07 return (K, c, cd)	14 return (K', cd')

KeyGen()	Decaps($sk = (sk', z), c$)
15 $(pk, sk') \leftarrow \text{PKE}_C.\text{KeyGen}()$	24 $m' \leftarrow \text{Dec}(sk', c)$
16 $z \leftarrow \mathcal{M}$	25 (c', cd')
17 $sk \leftarrow (sk', z)$	$\leftarrow \text{Enc}_C(pk, m'; G(m'))$
18 return (pk, sk)	28 $K' \leftarrow H(m', \text{ID}(pk))$
<u>Encaps(pk)</u>	29 $K'' \leftarrow H(z, c)$
19 $m \leftarrow \mathcal{M}$	30 $cd' \leftarrow J(m', cd')$
20 $(c, cd) \leftarrow \text{Enc}_C(pk, m'; G(m))$	31 $cd'' \leftarrow J(z, c)$
21 $K \leftarrow H(m, \text{ID}(pk))$	32 if $m' = \perp$ or $c' \neq c$ then
22 $cd \leftarrow J(m, cd)$	33 return (K'', cd'')
23 return (K, c, cd)	34 return (K', cd')

Public-code variants mask the code with an extra hash so that it does not leak the rejection event.

Figures from Glabush, Günther, Hale, Van Oorschot, Stebila. Augmenting post-quantum KEMs for verifiable decapsulation with rejection hiding. Under review, 2026.

Rejection hiding definition

$\text{Exp}_{\text{KEM}}^{\text{RH-ATK}}(\mathcal{A})$	$\mathcal{O}_G(m)$
01 $(pk, sk) \leftarrow \text{KeyGen}()$	11 $r \leftarrow G(m)$
02 $b \leftarrow \{0, 1\}$	12 $(c, K) \leftarrow \text{Encaps}^G(pk; m)$
03 $(K_0, c_0) \leftarrow \text{Encaps}^G(pk)$	13 $\mathcal{L}_G \leftarrow \mathcal{L}_G \cup \{c\}$
04 $b' \leftarrow \mathcal{A}^{\text{Chall}, \mathcal{O}_G, \mathcal{O}_{\text{Decaps}}}(pk, c_0)$	14 return r
05 if $c_1 \in \mathcal{L}_G \cup \mathcal{L}_D$:	$\mathcal{O}_{\text{Decaps}}(c \neq c_0)$
06 $b' \leftarrow 0$ // adversary loses	15 $K' \leftarrow \text{Decaps}^G(sk, c)$
07 return $\llbracket b' = b \rrbracket$	16 $\mathcal{L}_D \leftarrow \mathcal{L}_D \cup \{c\}$
<u>Chall($c \neq c_0$) // queried exactly once</u>	17 return K'
08 $c_1 \leftarrow c$	
09 $K' \leftarrow \text{Decaps}^G(sk, c_b)$	
10 return K'	

- Adversary must distinguish whether decapsulation of an invalid ciphertext *rejected* – given the full output (shared secret, and code if public).
- Theorem: HHK's original $\text{FO}^{\mathcal{Y}}$ is RH-CCA in the ROM, with

$$\text{Adv}_{\text{KEM}}^{\text{RH-CCA}} \leq \text{Adv}_{\text{KEM}}^{\text{IND-CCA}} + \frac{q_H}{|\mathcal{M}|} + 2^{-\gamma} + \delta$$