

How Multi-Recipient KEMs can help the Deployment of Post-Quantum Cryptography

Joël Alwen
AWS

Matthew Campagna
AWS

Dominik Hartmann
AWS

Shuichi Katsumata
PQShield & AIST

Eike Kiltz
Ruhr University
Bochum

Jake Massimo
AWS

Marta Mularczyk
AWS

Guillermo Pascual-Perez
ISTA

Thomas Prest
PQShield

Peter Schwabe
MPI & Radboud
University

Fifth PQC Standardization Conference

- Encapsulating K to 1 party using Kyber: **768 bytes**
- Encapsulating K to 100 parties using Kyber: **76 800 bytes**
- Encapsulating K to 100 parties using a “multi-recipient Kyber”:
5 504 bytes

How do we gain this factor 14?


Multi-Recipient KEMs



Main question

How efficiently can we share a session key K between $(N + 1)$ users?

- **Naive solution with El Gamal:**
 - Send $(g^{r_i}, pk_i^{r_i} \cdot K)$ for each user i
- **Variant by Kurosawa, PKC 2002:**
 - Send $(g^r, pk_1^r \cdot K, \dots, pk_N^r \cdot K)$
 - Asymptotically, saves a factor 2



Definition. In a decomposable encryption scheme, a ciphertext can be decomposed in key-dependent and key-independent parts:

$$\text{Enc}(pk_i, \text{msg}) = \underbrace{\text{Enc}^{\text{ind}}(r_0)}_{\text{ctxt}_0} \underbrace{\text{Enc}^{\text{dep}}(pk_i, \text{msg}, r_0, r_i)}_{\widehat{\text{ctxt}}_i}$$
The diagram illustrates the decomposition of a ciphertext. On the left, a dark blue rounded rectangle contains the text "Enc(pk_i, msg)". To its right is an equals sign. Further right are two more dark blue rounded rectangles. The first is labeled "ctxt_0" and is bracketed above by the text "Enc^{ind}(r_0)". The second is labeled with a wide-hat over "ctxt_i" and is bracketed above by the text "Enc^{dep}(pk_i, msg, r_0, r_i)".

Definition. In a decomposable encryption scheme, a ciphertext can be decomposed in key-dependent and key-independent parts:

$$\text{Enc}(pk_j, \text{msg}) = \underbrace{\text{ctxt}_0}_{\text{Enc}^{\text{ind}}(r_0)} \underbrace{\widehat{\text{ctxt}}_j}_{\text{Enc}^{\text{dep}}(pk_j, \text{msg}, r_0, r_j)}$$

El Gamal is decomposable. Let a ciphertext $\text{ctxt} = (g^r, pk_1^r \cdot \text{msg})$ with $pk_1 = g^{sk_1}$.

- 1 $\text{ctxt}_0 = g^r$.
- 2 $\widehat{\text{ctxt}}_1 = pk_1^r \cdot \text{msg}$.

A ciphertext with N recipients will be $\text{ctxt} \rightarrow (\text{ctxt}_0, \widehat{\text{ctxt}}_1, \dots, \widehat{\text{ctxt}}_N)$.
Key generation and decryption remain the same.

Definition. In a decomposable encryption scheme, a ciphertext can be decomposed in key-dependent and key-independent parts:

$$\text{Enc}(pk_j, \text{msg}) = \underbrace{\text{Enc}^{\text{ind}}(r_0)}_{\text{ctxt}_0} \underbrace{\text{Enc}^{\text{dep}}(pk_j, \text{msg}, r_0, r_j)}_{\widehat{\text{ctxt}}_j}$$

El Gamal is decomposable. Let a ciphertext $\text{ctxt} = (g^r, pk_1^r \cdot \text{msg})$ with $pk_1 = g^{sk_1}$.

- 1 $\text{ctxt}_0 = g^r$.
- 2 $\widehat{\text{ctxt}}_1 = pk_1^r \cdot \text{msg}$.

A ciphertext with N recipients will be $\text{ctxt} = (\text{ctxt}_0, \widehat{\text{ctxt}}_1, \dots, \widehat{\text{ctxt}}_N)$.
Key generation and decryption remain the same.

Questions:

- 1 What about CCA security?
 - ✓ $(\exists \text{ decomposable IND-CPA mPKE}) \xrightarrow{F-O} (\exists \text{ decomposable IND-CCA mKEM})$.
- 2 Is Kyber securely decomposable?

mKyber: a
Kyber-based
mKEM



Keygen ()

- 1 Sample \mathbf{A} and short \mathbf{s}, \mathbf{e}
- 2 $\mathbf{b} \leftarrow \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$
- 3 $\text{dk} := (\mathbf{s}, \mathbf{E}), \text{ek} := \mathbf{b}$

Enc(ek, msg)

- 1 Sample short row vectors $\mathbf{r}, \mathbf{e}', \mathbf{e}''$
- 2 $\mathbf{u} \leftarrow \mathbf{r} \cdot \mathbf{A} + \mathbf{e}'$
- 3 $\mathbf{v} \leftarrow \mathbf{r} \cdot \mathbf{b} + \mathbf{e}'' + \text{Encode}(\text{msg})$
- 4 $\text{ctxt} := (\mathbf{u}, \mathbf{v})$

Dec(dk, ctxt)

- 1 $\text{msg} \leftarrow \text{Decode}(\mathbf{v} - \mathbf{u} \cdot \mathbf{S})$

This construction is decomposable:

- Use the same \mathbf{A} for all public keys.
- \mathbf{u} is then independent of e_k and msg .

Enc(e_k \mathbf{b} , msg)

- 1 Sample short matrices \mathbf{r} , \mathbf{e}' , \mathbf{e}''
- 2 $\mathbf{u} \leftarrow \mathbf{r}\mathbf{A} + \mathbf{e}'$
- 3 $\mathbf{v} \leftarrow \mathbf{r}\mathbf{b} + \mathbf{e}'' + \text{Encode}(\text{msg})$
- 4 $\text{ctxt} := (\mathbf{u}, \mathbf{v})$

This construction is decomposable:

- Use the same \mathbf{A} for all public keys.
- \mathbf{u} is then independent of ek and msg .

Enc(ek , \mathbf{b} , msg)

- 1 Sample short matrices $\mathbf{r}, \mathbf{e}', \mathbf{e}''$
- 2 $\mathbf{u} \leftarrow \mathbf{r}\mathbf{A} + \mathbf{e}'$
- 3 $\mathbf{v} \leftarrow \mathbf{r}\mathbf{b} + \mathbf{e}'' + \text{Encode}(msg)$
- 4 $ctxt := (\mathbf{u}, \mathbf{v})$



MultiEnc($\{ek_1, \dots, ek_N\}, msg$)

- 1 Sample short matrices \mathbf{r}, \mathbf{e}'
- 2 $\mathbf{u} \leftarrow \mathbf{r}\mathbf{A} + \mathbf{e}'$
- 3 For $i = 1, \dots, N$:
 - 1 Sample a short matrix \mathbf{e}_i''
 - 2 $\mathbf{v}_i \leftarrow \mathbf{r}\mathbf{b}_i + \mathbf{e}_i'' + \text{Encode}(msg)$
- 4 $ctxt := (\mathbf{u}, \mathbf{v}_1, \dots, \mathbf{v}_N)$

This construction is decomposable:

- Use the same \mathbf{A} for all public keys.
- \mathbf{u} is then independent of ek and msg .

$Enc(ek, \mathbf{b}, msg)$

- 1 Sample short matrices $\mathbf{r}, \mathbf{e}', \mathbf{e}''$
- 2 $\mathbf{u} \leftarrow \mathbf{r}\mathbf{A} + \mathbf{e}'$
- 3 $\mathbf{v} \leftarrow \mathbf{r}\mathbf{b} + \mathbf{e}'' + \text{Encode}(msg)$
- 4 $ctxt := (\mathbf{u}, \mathbf{v})$



$MultiEnc(\{ek_1, \dots, ek_N\}, msg)$

- 1 Sample short matrices \mathbf{r}, \mathbf{e}'
- 2 $\mathbf{u} \leftarrow \mathbf{r}\mathbf{A} + \mathbf{e}'$
- 3 For $i = 1, \dots, N$:
 - 1 Sample a short matrix \mathbf{e}_i''
 - 2 $\mathbf{v}_i \leftarrow \mathbf{r}\mathbf{b}_i + \mathbf{e}_i'' + \text{Encode}(msg)$
- 4 $ctxt := (\mathbf{u}, \mathbf{v}_1, \dots, \mathbf{v}_N)$

Are we done? No!

- 1 Security?
- 2 Efficiency?

What assumptions do we rely on?

	Kyber	mKyber
Public key security	MLWE, $O(1)$ samples	MLWE, $O(1)$ samples
Ciphertext security	MLWE, $O(1)$ samples	MLWE, $O(N)$ samples

Which attacks are relevant against MLWE?

	Primal (Lattice)	Dual (Lattice)	Arora-Ge (Algebraic)	BKW (Combinatorial)
$O(1)$ samples	✓	✓	-	-
$O(N)$ samples	✓	✓	✓	✓

What assumptions do we rely on?

	Kyber	mKyber
Public key security	MLWE, $O(1)$ samples	MLWE, $O(1)$ samples
Ciphertext security	MLWE, $O(1)$ samples	MLWE, $O(N)$ samples

Which attacks are relevant against MLWE?

	Primal (Lattice)	Dual (Lattice)	Arora-Ge (Algebraic)	BKW (Combinatorial)
$O(1)$ samples	✓	✓	-	-
$O(N)$ samples	✓	✓	✓	✓




Are we in trouble? **No.**

✓ Bit dropping on the v_i makes Arora-Ge + BKW hard to the point of irrelevance

	Parameters								Sizes in bytes		
	q	n	k	η_1	η_2	d_u	d_v	$ \text{msg} $	$ \text{ek} $	$ \text{u} $	$ \text{v} $
Kyber-512	3329	256	2	3	2	10	4	32	800	640	128
mKyber-512	3329	256	2	3	2	11	3	16	768	704	48

	Parameters								Sizes in bytes		
	q	n	k	η_1	η_2	d_u	d_v	$ \text{msg} $	$ \text{ek} $	$ \text{u} $	$ \text{v} $
Kyber-512	3329	256	2	3	2	10	4	32	800	640	128
mKyber-512	3329	256	2	3	2	11	3	16	768	704	48

Not covered in this talk (see paper):

-  We can achieve IND-CCA security
-  We can upgrade to adaptive security by doubling the ciphertext size (amKyber)
-  Parameter selection differs from the KEM setting

Application 1: Broadcast



One sender sends the same keying material K to N parties

- Example application: state synchronisation in HSM fleet
- Perfect fit for mKEM!
- Also slightly simpler than naive solution (no DEM)

One sender sends the same keying material K to N parties

- Example application: state synchronisation in HSM fleet
- Perfect fit for mKEM!
- Also slightly simpler than naive solution (no DEM)

Example:

😊 1 Kyber ciphertext:



One sender sends the same keying material K to N parties

- Example application: state synchronisation in HSM fleet
- Perfect fit for mKEM!
- Also slightly simpler than naive solution (no DEM)

Example:

😊 1 Kyber ciphertext:



👤 N Kyber ciphertexts:



One sender sends the same keying material K to N parties

- Example application: state synchronisation in HSM fleet
- Perfect fit for mKEM!
- Also slightly simpler than naive solution (no DEM)

Example:

😊 1 Kyber ciphertext:




👤 N Kyber ciphertexts:



😊 1 mKyber ciphertext for N parties:

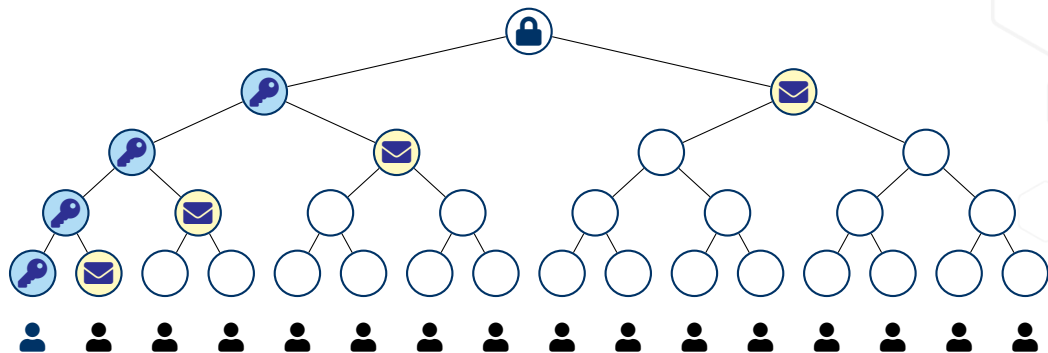


Application 2: MLS



The N users are arranged as the leaves of a (binary) tree

Tree invariant: (*user knows the private key of a node*) \Leftrightarrow (*node is in the path of user*)

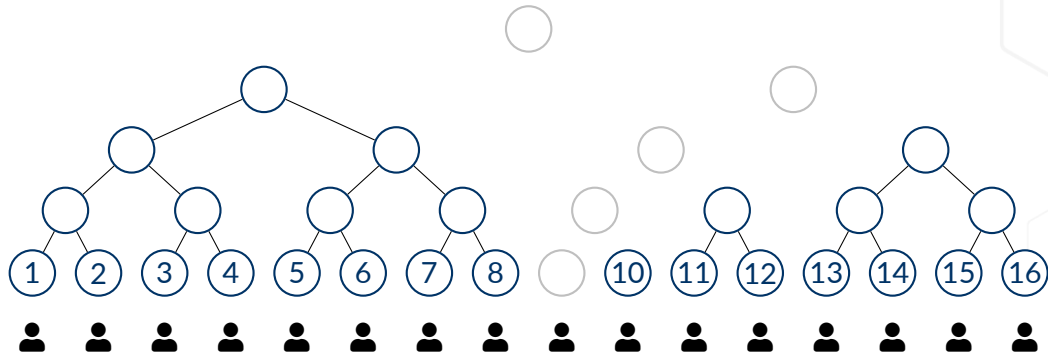


The N users are arranged as the leaves of a (binary) tree

Tree invariant: (*user knows the private key of a node*) \Leftrightarrow (*node is in the path of user*)

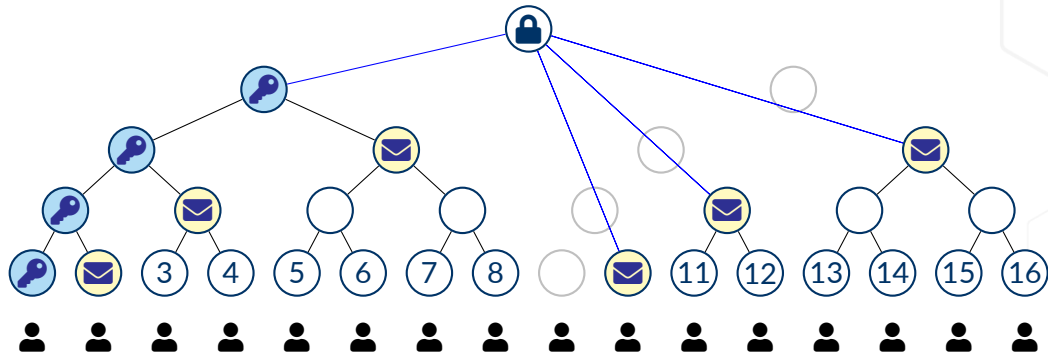
📶 Users routinely update their key material and broadcast:

- All $\lceil \log N \rceil$ encryption keys (🔑) in their direct path
- All $\geq \lceil \log N \rceil$ ciphertexts (✉) in their co-path
- 2 signatures (📝) - one for encryption keys, one for ciphertexts




When users are removed, their keys are removed for security.

→ This changes the topology of the tree




When users are removed, their keys are removed for security.

- This changes the topology of the tree
- This increases the number of ciphertext sent (here, 4 → 6)




Suppose we replace the binary tree by a star/flat tree:




Suppose we replace the binary tree by a star/flat tree:

- The number of ciphertexts become $O(N)$, but we can compress this using mKEM!



Suppose we replace the binary tree by a star/flat tree:

- The number of ciphertexts become $O(N)$, but we can compress this using mKEM!
- In addition, we can exploit the decomposability and have each user only download a portion $O(1)$ of the ciphertext






Suppose we replace the binary tree by a star/flat tree:

- The number of ciphertexts become $O(N)$, but we can compress this using mKEM!
- In addition, we can exploit the decomposability and have each user only download a portion $O(1)$ of the ciphertext

For more details: *More Efficient Protocols for Post-Quantum Secure Messaging*, RWC 2024. <https://www.youtube.com/watch?v=0hCPbu1wrhg>

Conclusion



-  mKEMs are a **simple and powerful tool** for scalable deployment of PQC
-  Many potential applications
-  We believe **standardizing mKEMs** would be useful

Questions?

