

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

Joël Alwen¹, Matthew Campagna¹, Dominik Hartmann², Shuichi Katsumata^{3,4}, Eike Kiltz², Jake Massimo¹, Marta Mularczyk¹, Guillermo Pascual-Perez⁵, Thomas Prest³, and Peter Schwabe⁶

¹Amazon Web Services

²Ruhr University Bochum

³PQShield

⁴AIST

⁵Institute of Science and Technology Austria (ISTA)

⁶Max Planck Institute for Security and Privacy, Radboud University Nijmegen

Abstract. The main purpose of this work is to raise awareness about a primitive that can provide large efficiency gains in post-quantum cryptography: multi-recipient KEMs, or mKEMs. In a nutshell, when encapsulating a key to N parties, an mKEM generates a single ciphertext that can be decapsulated by all parties. The size of an mKEM ciphertext can be significantly smaller than the sum of the sizes of N KEM ciphertexts. Moreover, individual receivers only need a small part of the mKEM ciphertext to run decapsulation.

We then propose mKyber, a very compact mKEM based on Kyber. Asymptotically, the size of an mKyber multi-recipient ciphertext is 16 times smaller than the sum of the sizes of N Kyber ciphertexts. The algorithmic description and parameters of mKyber and Kyber are very similar, which facilitates the re-use of existing security analyses, implementations, and formal verification tools that have been developed for Kyber.

Finally, we showcase some selected applications. mKEMs can be used to greatly reduce the bandwidth cost of the group key agreement protocol underlying the Message Layer Security (MLS) secure group messaging standard. Reducing bandwidth is one of the primary design considerations for MLS. More fundamentally, mKEMs reduce the cost of broadcasting private information to groups of recipients (e.g. a fleet of Cloud Hardware Security Modules).

1 Introduction

In July, 2022, NIST selected their first post-quantum standards for key agreement and (stateless) signature: the key encapsulation mechanism (KEM) Kyber [SAB⁺22], and the signature schemes Dilithium [LDK⁺22], SPHINCS⁺ [HBD⁺22] and Falcon [PFH⁺22]. While this will greatly speed up the transition of existing systems to post-quantum cryptography (PQC), some challenges will still need to be addressed in the process.

One of the main challenges in this transition process is the overhead in communication cost. For 128 bits of classical security, the size of an ECDH public key is 32 bytes, whereas the size of a Kyber ciphertext is 768 bytes, which is 24 times larger. This means that protocols that make an extensive use of key exchange or key encapsulation will require more bandwidth when migrating to PQC; such protocols include the IETF standard MLS [BBR⁺23], or broadcast protocols. This additional cost may require to scale up the bandwidth capabilities of the systems deploying these protocols, a cost that not all end users will be able to shoulder.

Multi-recipient KEMs (mKEMs). Fortunately, a simple primitive can help to address some of the scalability challenges faced by PQC: multi-recipient KEMs, or mKEMs. When encapsulating one key K to N distinct encapsulation keys $(ek_i)_{i \in [N]}$, the straightforward solution is to send N distinct ciphertexts ct_i – one for each ek_i . With a mKEM, one

may instead generate a single ciphertext \vec{c} that can be decrypted by all recipients. For a formal definition, see Section 2.

What makes mKEMs appealing for PQC is the massive efficiency gains they can provide. Recent works [KKPP20, HKP⁺21, AHK⁺23] have shown that lattice-based mKEMs can be asymptotically (in N) one or two orders of magnitude more compact than the use of N KEMs in parallel.

mKyber: a Kyber-based mKEM. We demonstrate the potential of mKEMs by proposing mKyber, a mKEM construction based on Kyber. When encrypting the same message to a large number of recipients, mKyber can be up to 16 times more compact than Kyber, with an amortized cost of 48 bytes per recipient compared to 768 bytes per recipient for Kyber. The algorithmic description and parameters of mKyber are very similar to those of Kyber. From an implementation perspective, it means that existing implementations of Kyber can be easily repurposed to implement mKyber, even in the case of highly optimized and/or platform-specific implementations. From a security evaluation perspective, security proof techniques and formal verification tools can be adapted. See Section 3 for more details.

amKyber: an even more secure mKEM. mKyber already enjoys the standard notion of IND-CCA security. However, in some settings a stronger notion of IND-CCA security *with adaptive corruptions* may be desired. Roughly, the latter requires IND-CCA security against an attacker that can adaptively leak mKEM secret keys. This is particularly important in applications using long-term mKEM key pairs. We therefore present amKyber, adapted from [AHK⁺23], which satisfies the stronger security notion. The amKyber public keys are only 256 bits larger than those of mKyber and the ciphertexts are only twice larger, i.e. 8 times more compact than for Kyber. The algorithms of amKyber are built by adapting in a simple way the building blocks of mKyber: the IND-CPA secure encryption scheme and the FO transform. Therefore, amKyber is easy to implement given an mKyber implementation and security proofs and formal verification tools for mKyber can be adapted.

Application: MLS and its variants. An immediate application of (post-quantum) mKEMs is to reduce the bandwidth cost of post-quantum deployments of secure group messaging such as the Messaging Layer Security (MLS) protocol (IETF RFC 9420 [BBR⁺23]). Minimizing bandwidth has been a primary design goal for MLS in the interest of scalability and supporting devices connected over low bandwidth networks (e.g. 2G cell phone networks). MLS was also designed specifically to make post-quantum secure deployments as easy as possible.¹

The MLS use case is particularly significant as MLS is proving to be an unexpectedly versatile tool with use cases already reaching well beyond its original domain of human-to-human secure messaging. MLS's scalable key agreement functionality for dynamic groups means it will serve as a key component underpinning higher-level multi-party cryptographic protocols such as Group Private Set Intersection, Privacy Preserving Federated Learning and end-to-end encrypted real-time voice/video communication. Its efficiency is also making it an attractive tool for IoT fleet command, control and coordination.

We show that mKEMs can greatly improve the scalability of MLS. Events such as removal of users can affect the efficiency of (standard) MLS, and mKEMs provide a simple way to mitigate this. In addition, the flexibility of mKEMs allows to come up with optimized variants of MLS. See Section 4 for more details.

Application: broadcast. The private broadcast functionality of mKEMs is a fundamental one. As such, post-quantum mKEMs have applications outside of MLS. For example, mKEMs allow for more efficient synchronization of fleets of (cloud based) Hardware Security Modules (HSMs). Such fleets underpin the security of major cloud providers. Privately synchronizing secret states across HSM fleets is indispensable for ensuring reliable operations and preventing loss due to hardware failure. See Section 5 for more details.

¹For example, this consideration was an important driver for the MLS working group's switch from the initial ART sub-protocol relying on Diffie-Hellman based Non-Interactive Key Agreement to the TreeKEM design, which instead made use of generic KEMs.

2 What is an mKEM?

2.1 Syntax

A multi-recipient key-encapsulation mechanism (mKEM) allows encapsulating a single key for multiple recipients. It consists of the following algorithms.

Parameter Generation: $\text{mKEM.Setup}() \rightarrow \text{pp}$ returns a fresh public parameter pp .

Key Generation: The key generation algorithm $\text{mKEM.Keygen}(\text{pp}) \rightarrow (\text{ek}, \text{dk})$ takes as input a public parameter pp and returns a fresh public/secret key pair (ek, dk) .

Encapsulation: The (multi-recipient) encapsulation algorithm $\text{mKEM.Encap}(\text{pp}, (\text{ek}_i)_{i \in [N]}) \rightarrow (\vec{\text{ct}}, K)$ takes in a sequence (of any length $n > 0$) of public keys and outputs a (multi-recipient) ciphertext $\vec{\text{ct}}$ and an encapsulated key K .

Extract: The deterministic algorithm $\text{mKEM.Ext}(\text{pp}, i, \vec{\text{ct}}) \rightarrow \text{ct}_i$ takes as input a position index i and a multi-recipient ciphertext $\vec{\text{ct}}$ and returns an individual ciphertext ct_i for the i -th recipient.

Decapsulation: The decapsulation algorithm $\text{mKEM.Decaps}(\text{pp}, \text{dk}_i, \text{ct}_i) \rightarrow K/\perp$ takes as input a secret key dk_i and an individual ciphertext ct_i . If decapsulation succeeds it returns the encapsulated key K (else \perp).

If the size of the individual ciphertext ct_i output by mKEM.Ext is independent of N , then we say that the mKEM is *efficiently decomposable*.

2.2 Security

We now define IND-CCA security for mKEMs, through what is a straightforward adaptation of the standard KEM security definition, with the difference that the adversary now gets challenged on a set of keys of its choosing, instead of a single public key. Our definition distinguishes between two different flavors, depending on whether (adaptive) corruptions by the adversary are allowed. Even though no specific attacks exploiting adaptive corruptions are known, the distinction is meaningful: known proof techniques for the weaker security notion fail when corruptions are allowed due to the so-called commitment problem. This is discussed further in Section 2.3.

The security experiment is described in Fig. 1. Roughly, it illustrates a game between a challenger and a stateful adversary \mathcal{A} that runs in two parts. The first one, \mathcal{A}_1 , gets a list of N public keys of an mKEM, of which it selects a subset on which it wants to be challenged. Then, \mathcal{A}_2 is given a ciphertext vector $\vec{\text{ct}}^*$ and a key K^* , and must decide whether K^* is a random key, independent of $\vec{\text{ct}}^*$, or $\vec{\text{ct}}^*$ is an encapsulation of K^* under the subset of keys chosen by \mathcal{A}_1 . In order to succeed, the adversary can make use of a state shared between both of its parts, as well as a decryption oracle, which allows it to query for decryptions of any ciphertext under any key, provided the query is not part of the challenge. For the stronger security notion where adaptive corruptions are allowed, \mathcal{A} is additionally given access to a corruption oracle, which simply returns the secret key of the queried key-pair. Finally, the adversary wins if it guesses correctly without having corrupted any of the keys from the challenge set.

Definition 1. Let mKEM be an mKEM scheme and N an integer. We say that mKEM is IND-CCA secure if, for all adversaries \mathcal{A} running in polynomial time,

$$\Pr[\text{Exp}_{\text{mKEM}, N, 1}^{\text{IND-CCA}}(\mathcal{A}) \rightarrow 1] - \Pr[\text{Exp}_{\text{mKEM}, N, 1}^{\text{IND-CCA}}(\mathcal{A}) \rightarrow 0]$$

is negligible in the security parameter.

2.3 Security with Adaptive Corruptions

Adaptive corruptions allow an adversary to decide, on the fly, what to corrupt depending on its full view (e.g. including public parameters, keys and ciphertexts). This makes it a strong, yet realistic type of adversary which is why various regular KEM and PKE schemes have been the subject of security analyses in adaptive corruptions models. For example,

<hr/> Alg. 1 Experiment $\text{Exp}_{\text{mKEM}, N, b}^{\text{IND-CCA}}(\mathcal{A})$ <hr/> 1: $(\mathcal{A}_1, \mathcal{A}_2) \leftarrow \mathcal{A}$ 2: $\text{pp} \leftarrow \text{mKEM.Setup}()$ 3: for $i \in [N]$ do 4: $(\text{ek}_i, \text{dk}_i) \leftarrow \text{mKEM.Keygen}(\text{pp})$ 5: $C \leftarrow \emptyset$ 6: $((i_j^* \in [N])_{j \in [n]}, st) \leftarrow \mathcal{A}_1^{\text{O}}(\text{pp}, (\text{ek}_i)_{i \in [N]})$ 7: $(\vec{\text{ct}}^*, K_0^*) \leftarrow \text{mKEM.Encap}(\text{pp}, (\text{ek}_{i_j^*})_{j \in [n]})$ 8: $K_1^* \leftarrow \{0, 1\}^\kappa$ 9: $b' \leftarrow \mathcal{A}_2^{\text{O}}(st, \vec{\text{ct}}^*, K_b^*)$ 10: return $b = b' \wedge \{i_j^* \mid j \in [n]\} \cap C = \emptyset$ <hr/>	<hr/> Alg. 2 Oracle $\text{Dec}(i \in [N], \text{ct})$ <hr/> 1: if $\text{ct}^* \neq \perp$ 2: if $\exists j \ i_j^* = i \wedge \text{ct} = \text{mKEM.Ext}(\text{pp}, j, \text{ct}^*)$ 3: return \perp 4: return $\text{mKEM.Decaps}(\text{pp}, \text{dk}_i, \text{ct})$ <hr/> <hr/> Alg. 3 Oracle $\text{Cor}(i \in [N])$ <hr/> 1: $C \leftarrow C \cup \{i\}$ 2: return dk_i <hr/>
--	---

Figure 1: IND-CCA security experiments for mKEM. With (adaptive) corruptions (IND-CCA^{a-mu}), the oracle set is $\text{O} := \{\text{Dec}, \text{Cor}\}$. Without corruptions (IND-CCA^{mu}), the oracle set is $\text{O} := \{\text{Dec}\}$.

their security is often analyzed in the multi-user setting, e.g. [BBM00, ABH⁺21]. Here, the IND-CCA adversary chooses which out of N (independently generated) key pairs to challenge. Thus we can define two multi-user security notions for an mKEM: the standard *non-adaptive* notion IND-CCA^{mu} in which the adversary decides prior to receiving any inputs which keys it will corrupt and IND-CCA^{a-mu}, the strictly stronger *adaptive* notion in which the adversary can decide on the fly as the execution progresses which (non-challenge) secret keys to corrupt.² We depict both in Figure 1.

Security proofs of mKEMs and KEMs typically consider IND-CCA^{mu}, e.g. [BBM00, ABH⁺21, Kur02, BBS03, BBKS07, PPS14, KKPP20]. For Diffie-Hellman (DH) based schemes this allows proving tighter security bounds, i.e., the security loss of the reduction to the DH assumption does not depend on N , see e.g. in [ABH⁺21] (HPKE) or in [Kur02] (a DHIES-based mKEM). Note that this affects the choice of concrete parameters for the scheme. The above proof techniques (which all use random self-reducibility) *cannot* be used to prove tight IND-CCA^{a-mu} security. To the best of our knowledge, all reductions to date proving IND-CCA^{a-mu} of KEMs or mKEMs lose either the factor N (using the standard hybrid argument) or at least the number of corrupted keys [AHKM22].

For the case of mKyber we do not know of any proof of IND-CCA^{a-mu} where the reduction has a non-exponential security loss in N (resulting from guessing the set of corrupted keys). We stress that, we believe it to be very unlikely that there is an actual attack on mKyber exploiting adaptivity in the multi-user setting. Instead, our current proof techniques do not work here for a technical (i.e. the so-called “commitment”) problem. It is possible that new proof techniques could be used in the future to prove IND-CCA^{a-mu} of mKyber.

We also note that in applications where key pairs live for very short time periods, which makes corruptions unlikely, IND-CCA^{mu} can be sufficient. This is the case e.g. for TLS. In contrast, for applications like MLS we can expect keys to live for significantly longer periods of time, especially in large groups.

3 Kyber-based mKEMs

This section defines two mKEM constructions: mKyber in Section 3.1 and amKyber in Section 3.2. The latter is secure with (adaptive) corruptions but slightly less efficient. We also

²Outside this section we use the term IND-CCA to refer to either, for simplicity. Which one is meant will be clear from the context.

propose in Section 3.3 concrete parameters for `mKyber` and `amKyber`. Assuming there are many recipients, instantiating (a)`mKyber` with these parameters results in shorter ciphertexts over instantiating them with the parameters of `Kyber` (which would be secure, though).

3.1 The `mKyber` Construction

Simplifications. For conciseness, the description of `Kyber` in Fig. 2 is heavily simplified compared to the one in [SAB⁺22]. For example:

- Some calls to hash functions or PRFs are omitted or merged together;
- We assume that \mathbf{A} is part of `ek`. For compactness, in practice `ek` instead contains a seed `seedA`, which can be extended to \mathbf{A} by passing it into a XOF: $\mathbf{A} := \text{XOF}(\text{seed}_{\mathbf{A}})$.
- We simplify details relative to the bit representation of mathematical objects. We also omit transitions between the NTT representation and the coefficient representation.

Making \mathbf{A} part of the public parameters. In MLWE-based KEMs such as `Kyber`, the matrix \mathbf{A} is different for each keypair. For `mKyber`, the matrix \mathbf{A} is the same for all keypairs in order to enable the benefits of mKEMs. See also [KKPP20].

Syntax. As explained in (previous section), the syntax of mKEMs is different than for regular KEMs. The encapsulation procedure (Algorithm 7) now takes N encapsulation keys instead of one as in a regular KEM (Algorithm 6).

Also note that for mKEMs, the encapsulation procedure outputs a multi-recipient ciphertext $\vec{\text{ct}} := (\mathbf{u}, (v_i)_{i \in [N]})$, but a user with the decryption key `dki` only needs the partial ciphertext $\text{ct}_i = (\mathbf{u}, v_i)$ as input to its decapsulation procedure (Algorithm 9).

Decomposable CCA transform. Our mKEM construction is efficiently decomposable, allowing to convert a multi-recipient ciphertext $\vec{\text{ct}} := (\mathbf{u}, (v_i)_{i \in [N]})$ of size $O(N)$ into a single-recipient ciphertext $\text{ct}_i = (\mathbf{u}, v_i)$ of size $O(1)$, where \mathbf{u} is independent of the encryption keys. Decomposability [KKPP20] can be exploited for further efficiency gains, see Section 4. In order to achieve decomposability in a CCA setting, we also require a CCA transform with a decomposable flavor. To this effect, we replace the Fujisaki-Okamoto transform [FO13, HHK17] used in `Kyber` by a decomposable variant introduced in [KKPP20].

3.2 The `amKyber` Construction Secure with Adaptive Corruptions

This section describes a variant of `mKyber` called `amKyber`, proposed in [AHK⁺23], which achieves the stronger adaptive security notion of IND-CCA^{a-mu} defined in Figure 1. We describe `amKyber` in Figs. 3 and 4.

Like `mKyber` (and `Kyber`), the `amKyber` mKEM is built by applying an FO transform to an underlying CPA secure mPKE (or PKE in the case of `Kyber`). As shown in [AHK⁺23], the FO transform preserves the (quantum and classic) adaptive security of the underlying mPKE resulting in an adaptively secure mKEM. Thus, en lieu of the mPKE underlying `mKyber` (which we shall call `mKyberPKE`), we use one called `amKyberPKE` which [AHK⁺23] shows to be an adaptively secure IND-CPA mPKE. `amKyber` is then obtained from `amKyberPKE` by applying the same FO as produces `mKyber` from `mKyberPKE`.

In more detail, `amKyberPKE` adapts to the mPKE setting, the ideas of [GW09] for building adaptively-secure broadcast encryption. At a high level, `amKyberPKE` runs two parallel instances of `mKyberPKE`. In particular, an `amKyberPKE` public key consists of two `mKyberPKE` public keys, \mathbf{b}_l (left) and \mathbf{b}_r (right) but the corresponding `amKyberPKE` secret key contains just one of the two matching `mKyberPKE` secret keys. Which one is chosen privately and at random during key generation. As such, we can use an alternative public key generation method for the other key pair which does not produce the corresponding secret key but therefore allows for a more compact representation of the resulting `amKyberPKE` public key. Namely, rather than encoding the two `mKyberPKE` public keys explicitly as $(\mathbf{b}_l, \mathbf{b}_r)$, they are instead encoded as the (much shorter) pair $(\mathbf{b}_l, \text{seed})$ where `seed` is a string of 256 bits. Together they define the second public key to be $\mathbf{b}_r := \text{HashToEk}(\text{seed}) - \mathbf{b}_l$ where `HashToEk`

<p>Alg. 4 Kyber.Keygen(pp)</p> <ol style="list-style-type: none"> 1: seed $\leftarrow \{0, 1\}^{256}$ 2: seed' $\leftarrow \{0, 1\}^{256}$ 3: $\mathbf{A}, \mathbf{s}, \mathbf{e} := \text{PRF}(\text{seed}')$ $\triangleright \mathbf{A} \in R_q^{k \times k}$ 4: $\mathbf{b} := \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ $\triangleright \mathbf{b}, \mathbf{s}, \mathbf{e} \in R_q^k$ 5: ek := (\mathbf{A}, \mathbf{b}) 6: dk := (s, ek, seed) 7: return ek, dk 	<p>Alg. 5 mKyber.Keygen(pp $\ni \mathbf{A}$)</p> <ol style="list-style-type: none"> 1: seed $\leftarrow \{0, 1\}^{256}$ 2: seed' $\leftarrow \{0, 1\}^{256}$ 3: $\mathbf{s}, \mathbf{e} := \text{PRF}(\text{seed}')$ 4: $\mathbf{b} := \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ 5: ek := \mathbf{b} 6: dk := (s, ek, seed) 7: return ek, dk
<p>Alg. 6 Kyber.Encap(ek = \mathbf{b})</p> <ol style="list-style-type: none"> 1: msg $\leftarrow \{0, 1\}^{256}$ 2: msg := H(msg) 3: (K', coin) := G(msg, H(ek)) 4: $\mathbf{r}, \mathbf{e}', \mathbf{e}'' := \text{PRF}(\text{coin})$ 5: $\mathbf{u} := \mathbf{r} \cdot \mathbf{A} + \mathbf{e}'$ $\triangleright \mathbf{u}, \mathbf{r}, \mathbf{e}' \in R_q^{1 \times k}$ 6: $\mathbf{v} := \mathbf{r} \cdot \mathbf{b} + \mathbf{e}'' + \text{Encode}(\text{msg})$ $\triangleright \mathbf{v}, \mathbf{e}'' \in R_q$ 7: $\bar{\mathbf{u}} := \text{Compress}(\mathbf{u}, d_u)$ 8: $\bar{\mathbf{v}} := \text{Compress}(\mathbf{v}, d_v)$ 9: return ct := ($\bar{\mathbf{u}}, \bar{\mathbf{v}}$), $K := \text{KDF}(K', \text{ct})$ 	<p>Alg. 7 mKyber.Encap($(\text{ek}_i = \mathbf{b}_i)_{i \in [N]}$)</p> <ol style="list-style-type: none"> 1: msg $\leftarrow \{0, 1\}^k$ 2: msg := H(msg) 3: $K' := \text{H}(\text{msg})$ 4: coin := $G_1(\text{msg})$ 5: $\mathbf{r}, \mathbf{e}' := \text{PRF}_1(\text{coin})$ 6: $\mathbf{u} := \mathbf{r} \cdot \mathbf{A} + \mathbf{e}'$ 7: $\bar{\mathbf{u}} := \text{Compress}(\mathbf{u}, d_u)$ 8: for $i \in [N]$ do 9: $\text{coin}_i := G_2(\text{ek}_i, \text{msg})$ 10: $\mathbf{e}''_i := \text{PRF}_2(\text{coin}_i)$ 11: $\mathbf{v}_i := \mathbf{r} \cdot \mathbf{b}_i + \mathbf{e}''_i + \text{Encode}(\text{msg})$ 12: $\bar{\mathbf{v}}_i := \text{Compress}(\mathbf{v}_i, d_v)$ 13: $\bar{\text{ct}} := (\bar{\mathbf{u}}, (\bar{\mathbf{v}}_i)_{i \in [N]})$ 14: return ct, $K := \text{KDF}(K', \text{ct})$
<p>Alg. 8 Kyber.Decaps(dk, ct = ($\bar{\mathbf{u}}, \bar{\mathbf{v}}$))</p> <ol style="list-style-type: none"> 1: ($\mathbf{s}, \mathbf{b}, \text{seed}$) \leftarrow dk 2: $\mathbf{u} := \text{Decompress}(\bar{\mathbf{u}}, d_u)$ 3: $\mathbf{v} := \text{Decompress}(\bar{\mathbf{v}}, d_v)$ 4: msg := Decode($\mathbf{v} - \mathbf{u} \cdot \mathbf{s}$) 5: ($K', \text{coin}$) := G(msg, H(ek)) 6: $\bar{\mathbf{r}}, \bar{\mathbf{e}}', \bar{\mathbf{e}}'' := \text{PRF}(\text{coin})$ 7: $\bar{\mathbf{u}} := \bar{\mathbf{r}} \cdot \mathbf{A} + \bar{\mathbf{e}}'$ 8: $\bar{\mathbf{v}} := \bar{\mathbf{r}} \cdot \mathbf{b} + \bar{\mathbf{e}}'' + \text{Encode}(\text{msg})$ 9: $\bar{\bar{\mathbf{u}}} := \text{Compress}(\bar{\mathbf{u}}, d_u)$ 10: $\bar{\bar{\mathbf{v}}} := \text{Compress}(\bar{\mathbf{v}}, d_v)$ 11: $\bar{\text{ct}} := \bar{\bar{\mathbf{u}}}, \bar{\bar{\mathbf{v}}}$ 12: if $\bar{\text{ct}} = \text{ct}$ 13: return $K := \text{KDF}(K', \text{ct})$ 14: else 15: return $K := \text{KDF}(\text{seed}, \text{ct})$ 	<p>Alg. 9 mKyber.Decaps($\text{dk}_i, \text{ct}_i = (\bar{\mathbf{u}}, \bar{\mathbf{v}}_i)$)</p> <ol style="list-style-type: none"> 1: ($\mathbf{s}_i, \mathbf{b}_i, \text{seed}_i$) \leftarrow dk 2: $\mathbf{u} := \text{Decompress}(\bar{\mathbf{u}}, d_u)$ 3: $\mathbf{v}_i := \text{Decompress}(\bar{\mathbf{v}}_i, d_v)$ 4: msg := Decode($\mathbf{v}_i - \mathbf{u} \cdot \mathbf{s}_i$) 5: $K' := \text{H}(\text{msg})$ 6: coin := $G_1(\text{msg})$ 7: $\mathbf{r}, \mathbf{e}' := \text{PRF}_1(\text{coin})$ 8: $\text{coin}_i := G_2(\text{ek}_i, \text{msg})$ 9: $\mathbf{e}''_i := \text{PRF}_2(\text{coin}_i)$ 10: $\mathbf{u} := \mathbf{r} \cdot \mathbf{A} + \mathbf{e}'$ 11: $\mathbf{v}_i := \mathbf{r} \cdot \mathbf{b}_i + \mathbf{e}''_i + \text{Encode}(\text{msg})$ 12: $\bar{\mathbf{u}} := \text{Compress}(\mathbf{u}, d_u)$ 13: $\bar{\mathbf{v}}_i := \text{Compress}(\mathbf{v}_i, d_v)$ 14: $\bar{\text{ct}}_i := (\bar{\mathbf{u}}, \bar{\mathbf{v}}_i)$ 15: if $\bar{\text{ct}}_i = \text{ct}_i$ 16: return $K := \text{KDF}(K', \text{ct}_i)$ 17: else 18: return $K := \text{KDF}(\text{seed}_i, \text{ct}_i)$

Re-encryption

Re-encryption

Figure 2: Side-by-side comparison of Kyber (Algorithms 4, 6 and 8) and mKyber (Algorithms 5, 7 and 9). The main algorithmic differences are highlighted, and we discuss them in Section 3.1.

is a public hash function mapping seeds to public keys. To generate an `amKyberPKE` key pair first choose a uniform random `seed`. Next flip a fair coin and, depending on the outcome, generate either (dk_l, \mathbf{b}_l) or (dk_r, \mathbf{b}_r) using `mKyberPKE` key generation algorithm. Finally, set remaining `mKyberPKE` public key such that $\mathbf{b}_l := \text{HashToEk}(\text{seed}) - \mathbf{b}_r$.

Alg. 10 `amKyber.Keygen(pp = A)`

```

1:  $(ek', dk') := \text{mKyber.Keygen}(pp)$ 
2:  $\text{seed} \leftarrow \{0, 1\}^{256}$ 
3:  $\text{swpEk} \leftarrow \{0, 1\}$ 
4: if  $\text{swpEk} = 1$ 
5:    $ek' := \text{HashToEk}(\text{seed}) - ek'$ 
6: return  $ek := (ek', \text{seed}), dk := (dk', \text{swpEk}, ek', \text{seed})$ 

```

Alg. 11 `mKyber.Encap` $((ek_i = \mathbf{b}_i)_{i \in [N]})$

```

1:  $\text{msg} \leftarrow \{0, 1\}^\kappa$ 
2:  $\text{msg} := H(\text{msg})$ 
3:  $K' := H(\text{msg})$ 
4:  $(u, *) := \text{CpaEncU}(\text{msg})$ 
5: for  $i \in [N]$  do
6:    $v_i := \text{CpaEncV}(\text{msg}, ek_i, G_2(\text{msg}, ek_i))$ 
7:  $\vec{ct} := (u, (v_i)_{i \in [N]})$ 
8: return  $\vec{ct}, K := \text{KDF}(K', ct)$ 

```

Alg. 12 `amKyber.Encap` $((ek_i)_{i \in [N]})$

```

1:  $\text{msg} \leftarrow \{0, 1\}^\kappa$ 
2:  $\text{msg} := H(\text{msg})$ 
3:  $K' := H(\text{msg})$ 
4:  $\text{msg}_l, \text{msg}_r := G_3(\text{msg})$ 
5:  $(u_l, *) := \text{CpaEncU}(\text{msg}_l)$ 
6:  $\mathbf{u} := (u_l, u_r)$ 
7:  $(u_r, *) := \text{CpaEncU}(\text{msg}_r)$ 
8: for  $i \in [N]$  do
9:    $(\mathbf{b}_l, \text{seed}) := ek_i$ 
10:   $\mathbf{b}_r := \text{HashToEk}(\text{seed}) - \mathbf{b}_l$ 
11:   $\text{swpC} := G_4(\text{msg}, ek_i) \in \{0, 1\}$ 
12:  if  $\text{swpC} = 1$ 
13:     $(\mathbf{b}_l, \mathbf{b}_r) \leftarrow (\mathbf{b}_r, \mathbf{b}_l)$ 
14:   $v_l := \text{CpaEncV}(\text{msg}, \mathbf{b}_l, G_2(\text{msg}_l, ek_i))$ 
15:   $v_r := \text{CpaEncV}(\text{msg}, \mathbf{b}_r, G_2(\text{msg}_r, ek_i))$ 
16:   $v_i := (v_l, v_r, \text{swpC})$ 
17:  $\vec{ct} := (u, (v_i)_{i \in [N]})$ 
18: return  $\vec{ct}, K := \text{KDF}(K', ct)$ 

```

Figure 3: Key generation of `amKyber` (Algorithm 10) as well as a side-by-side comparison of the encapsulation of `mKyber` (Algorithm 11) and `amKyber` (Algorithm 12). The main algorithmic differences are highlighted, and we discuss them in Section 3.2.

To encrypt `msg` to multiple recipients, `amKyberPKE` also runs `mKyberPKE`'s encryption twice: It first generates two vectors \mathbf{u} called \mathbf{u}_l and \mathbf{u}_r . Then, for each recipient, it flips a fair coin (using the random oracle on `msg` as required by the FO transform) to decide which of the recipient's keys, \mathbf{b}_l or \mathbf{b}_r , goes to the recipients left `mKyber` instance (i.e. with public key \mathbf{u}_l) and which goes to the right instance. To decrypt `msg`, the recipient in `amKyberPKE` runs `mKyberPKE` with the secret key it knows and the corresponding invocation indicated by the sender.

Security. `amKyber` uses the same parameters as `mKyber`. The adaptive security proof of `amKyber` can be found in [KKPP20, AHK⁺23]. In particular, [KKPP20] proves `mKyberPKE` to be (non-adaptively) IND-CPA secure mPKE. Next, in [AHK⁺23] proves their transformation of `mKyberPKE` into `amKyberPKE` results in an *adaptively* IND-CPA secure mPKE.

<hr/> <p>Alg. 13 mKyber.Decaps($dk_i, ct_i = (u, v_i)$)</p> <hr/> <pre> 1: $(s_i, b_i, seed_i, ek') \leftarrow dk$ 2: $u := Decompress(u, d_u)$ 3: $v_i := Decompress(v_i, d_v)$ 4: $msg := Decode(v_i - u \cdot s_i)$ 5: $K' := H(msg)$ 6: $(\bar{u}, *) := CpaEncU(msg)$ 7: $\bar{v} := CpaEncV(msg, b_i, G_2(msg, ek_i))$ 8: $\bar{ct}_i := (\bar{u}, \bar{v})$ 9: if $\bar{ct}_i = ct_i$ 10: return $K := KDF(K', ct_i)$ 11: else 12: return $K := KDF(seed_i, ct_i)$ </pre> <hr/>	<hr/> <p>Alg. 14 amKyber.Decaps($dk_i, ct_i = (u, v_i)$)</p> <hr/> <pre> 1: $(dk', swpEk, b_l, seed) := dk_i$ 2: $b_r := HashToEk(seed) - b_l$ 3: $(v_l, v_r, swpC) := v_i$ 4: $(u_l, u_r) := u_i$ 5: $ek_i := (b_l, seed)$ 6: if $swpEk \text{ XOR } swpC = 0$ 7: $side := l$ 8: else 9: $side := r$ 10: $(s, b, seed) := dk'$ 11: $u := Decompress(u_{side}, d_u)$ 12: $v := Decompress(v_{side}, d_v)$ 13: $msg := Decode(v - u \cdot s)$ 14: $K' := H(msg)$ 15: $msg_l, msg_r := G_3(msg)$ 16: for $side \in \{l, r\}$ do 17: $(\bar{u}_{side}, *) := CpaEncU(msg_{side})$ 18: $coin := G_2(msg_{side}, ek_i)$ 19: $\bar{v}_{side} := CpaEncV(msg, b_{side}, coin)$ 20: $swpC := G_4(msg, ek_i)$ 21: if $(\bar{u}_l, \bar{v}_l, \bar{u}_r, \bar{v}_r, swpC) = (u_l, v_l, u_r, v_r, swpC)$ 22: return $K := KDF(K', ct_i)$ 23: else 24: return $K := KDF(seed_i, ct_i)$ </pre> <hr/>
<hr/> <p>Alg. 15 CpaEncU(msg)</p> <hr/> <pre> 1: $r, e' := PRF_1(G_1(msg))$ 2: $u := r \cdot A + e'$ 3: $u := Compress(u, d_u)$ 4: return (u, r) </pre> <hr/>	<hr/> <p>Alg. 16 CpaEncV(msg, $ek_i = b_i, coin_i$)</p> <hr/> <pre> 1: $(*, r) := CpaEncU(msg)$ 2: $e''_i := PRF_2(coin_i)$ 3: $v_i := r \cdot b_i + e''_i + Encode(msg)$ 4: return $v_i := Compress(v_i, d_v)$ </pre> <hr/>

Figure 4: Side-by-side comparison of decapsulation of mKyber (Algorithm 11) and amKyber (Algorithm 12). The main algorithmic differences are highlighted, and we discuss them in Section 3.2.

The same paper also proves that the FO transform of [KKPP20] applied to `amKyberPKE` produces `amKyber`; an adaptively IND-CCA^{a-mu} secure mKEM. We note that all of the above security statements hold both against classical and quantum adversaries.

Together, these results provide strong evidence for the security of `amKyber` design.

Efficiency. An `amKyber` public key has only 256 bits in addition to an `mKyber` public key. An `amKyber` ciphertext is twice larger than an `mKyber` ciphertext. Note that asymptotically, `amKyber` is still 8 times smaller than the sum of sizes of N Kyber ciphertexts.

3.3 Parameters Selection

In Table 1, we propose concrete `mKyber` parameters targeting the NIST security level I (at least as hard as key-recovery on AES-128). As discussed in Section 1, the parameters are largely similar to those of `Kyber`; we only tweaked the parameters $(d_u, d_v, |\text{msg}|)$, as it allows to greatly decrease the size of $|\mathbf{v}|$. While `Kyber` and `mKyber` are largely similar, the mKEM setting impact both the efficiency analysis and the security analysis. We analyse both of them separately.

Table 1: Parameter sets of `Kyber512` and `mKyber512`.

	Parameters								Sizes in bytes		
	q	n	k	η_1	η_2	d_u	d_v	$ \text{msg} $	$ \text{ek} $	$ \mathbf{u} $	$ \mathbf{v} $
<code>Kyber512</code>	3329	256	2	3	2	10	4	32	800	640	128
<code>mKyber512</code>	3329	256	2	3	2	11	3	16	768	704	48

Efficiency. Let us note $|x|$ the size in bytes of an object. In the KEM regime, it is of interest for most applications to minimize the ciphertext size $|\text{ct}|$, the encapsulation key size $|\text{ek}|$, or some linear combination of the two. In the mKEM regime, a multi-recipient ciphertext is of the form $\vec{\text{ct}} := (\mathbf{u}, (\mathbf{v}_i)_{i \in [N]})$. Therefore there is a high incentive to minimize $|\mathbf{v}_i|$, since asymptotically $|\vec{\text{ct}}| \sim N \cdot |\mathbf{v}_i|$. There are a few tricks we can use to minimize $|\mathbf{v}_i|$.

1. **Shorter msg.** In `Kyber`, the message `msg` is 256 bits long across all parameter sets. However, if we note κ is the targeted bit-security level (i.e. $\kappa \in \{128, 192, 256\}$ for the NIST level I, III, V), it suffices to take a message of κ bits.
2. **Coefficient dropping.** When applying the previous idea, we only need to encode κ bits in a polynomial v_i of 256 coefficients. We only need to send κ coefficients of v_i instead of n coefficients. In effect, this divides $|\mathbf{v}_i|$ by factor $256/\kappa$, which is a factor two when $\kappa = 128$.
3. **Bit dropping.** It is customary to apply bit dropping on \mathbf{u} and \mathbf{v} . This is parametrized by d_u and d_v , which are the number of bits sent per coefficient of \mathbf{u} and \mathbf{v}_i , respectively. The bitsize of \mathbf{v}_i is therefore $\kappa \cdot d_v$.

Since our goal is to minimize $|\mathbf{v}_i|$, we reduce d_v from 4 to 3. This increases the decryption failure rate (DFR), so we increase d_u from 10 to 11 in order to keep the DFR low.

Putting these optimisations together, we manage to obtain $|\mathbf{v}_i| = 48$.

Security. Both `Kyber` and `mKyber` rely on MLWE for the security of the encapsulation key. For the ciphertext, due to the simultaneous presence of additive noise and rounding (via `Compress`), they both rely on a hybrid between MLWE and MLWR which we will call MLWER. The major difference is that a `mKyber` multi-ciphertext $\vec{\text{ct}} := (\mathbf{u}, (\mathbf{v}_i)_{i \in [N]})$ contains a number of MLWER samples that is affine in N . This new parameter has varying impacts on the existing methods to solve MLWE, MLWR and MLWER. There are roughly three families of such methods:

1. **Lattices.** Methods based on (primal or dual) lattice reduction are usually the most relevant for most parametrizations of lattice-based schemes. For these methods, having a large number of samples provide results little to no advantage.

2. **Algebraic.** Algebraic methods, such as Arora-Ge and its variants, rely on linearization. They require a large number of samples: $(k \cdot n)^D$, where D is the size of the support of the errors (both due to additive noise and to rounding). While the required number of samples is too large in the KEM regime, mKEM ciphertexts contain a large number of MLWER samples. The relevance of Arora-Ge therefore needs to be re-assessed.

Fortunately, the ciphertexts in mKyber are very noisy. Indeed, each v_i undergoes heavy bit dropping on each coefficient, see Table 1. This makes the required number of samples much larger than 2^κ , which rules out the Arora-Ge attack in its current form.

3. **Combinatorial.** Combinatorial methods, such as BKW, combine lattice attacks and guessing. Like Arora-Ge, BKW and its variants need a large number of samples and therefore need to be considered for mKyber. Fortunately, they are also very sensitive to the size of the errors' supports. Due to the heavy rounding performed on the v_i , BKW requires an intractable number of samples and is in effect inoperative against the parameters in Table 1.

A fully detailed discussion on the security analysis of mKEMs against known attacks can be found in [HKP⁺21, Appendix G]. The insights and trade-offs discussed in this section are summarized in Table 2.

Table 2: Impact of parameters on security and performance metrics. Terminology: \nearrow (resp. $=$, resp. \searrow) indicates that increasing this parameter has a positive (resp. essentially neutral, resp. negative) impact on the considered metric.

	Communication cost			Correctness	Ciphertext security		
	ek	u	v_i		Arora-Ge	BKW	Lattice
N	$=$	$=$	$=$	$=$	\searrow	\searrow	$=$
q	\searrow	$=$	$=$	\nearrow	$=$	$=$	\searrow
n	\searrow	\searrow	$=$	\searrow	\nearrow	\nearrow	\nearrow
k	\searrow	\searrow	$=$	\searrow	\nearrow	\nearrow	\nearrow
η_1	$=$	$=$	$=$	\searrow	\nearrow	\nearrow	\nearrow
η_2	$=$	$=$	$=$	\searrow	\nearrow	\nearrow	\nearrow
d_u	$=$	\searrow	$=$	\searrow	\nearrow	\nearrow	\nearrow
d_v	$=$	$=$	\searrow	\searrow	\nearrow	\nearrow	\nearrow

4 Application: MLS and its Variants

MLS and TreeKEM. Messaging Layer Security (MLS), is a protocol for end-to-end security. MLS has been standardized by IETF under RFC 9420 [BBR⁺23] in March 2023, and enjoys the support of several industry actors (Google, Amazon, Cisco, etc.).

The notion of continuous group key agreement (CGKA) was put forward to capture the notion of secure group management which lies at the core of secure group messaging and other protocols. The main function of a CGKA is to ensure the secure distribution of a shared key K inside a group of N users connected to an untrusted server. A CGKA must be able to support the addition or removal of users to the group, as well as properties such as forward secrecy and post-compromise security.

Inside MLS, the sub-protocol TreeKEM³ performs the function of a CGKA. In TreeKEM, the N users inside a group are positioned at the leaves of a binary tree called the ratchet tree, see Fig. 5a. To each node i of the ratchet tree is associated an encapsulation keypair (ek_i, dk_i) . The ratchet tree maintains the following tree invariant [BBR⁺23, §4.2]:


(TI) “The private key for a node in the tree is known to a member of the group if and only if the node’s subtree contains that member’s leaf.”

³While the term “TreeKEM” itself does not appear in the MLS RFC [BBR⁺23], it is implemented through the concept of “ratchet trees”. The term TreeKEM is commonly used as shorthand in research articles to refer to this particular part of MLS. We follow the same practice here.


A critical operation in MLS/TreeKEM consists of a user refreshing their cryptographic keys inside the ratchet tree. They do so via the following steps:

1. Generate a node secret s_0 for their leaf node;
2. For each node i in their direct path, i.e. the path from their leaf node to the root, starting from the leaf node, use the node secret s_i as follows:
 - (a) Pass s_i into a PRF to derive its parent node's secret $s_{\text{parent}(i)} = s_{i+1}$ (except when i is the root).
 - (b) Pass s_i through a PRF with a different input than in Item 2a and use the output to derive a KEM key pair (ek_i, dk_i) .
 - (c) Compute a ciphertext ct_i encrypting s_{i+1} under the encapsulation key of the sibling node of i using the KEM-DEM paradigm and concretely HPKE, defined in RFC 9180 [BBLW22]. (Except for the leaf node)
3. Broadcast a *commit message* containing all the (ek_i, ct_i) for i on the path of the user.


One can show that this preserves the tree invariant (TI). This method enjoys great communication complexity; a commit message contains $\lceil \log N \rceil$ public keys and as many ciphertexts.




(a) TreeKEM, complete ratchet tree



(b) TreeKEM after a user has been removed and their direct path blanked out. The three boxed ciphertexts encrypt the same value: the node secret of the root.



(c) m -ary TreeKEM [KKPP20], full ratchet tree



(d) Chained CmPKE [HKP+21]

Figure 5: Illustration of the refresh operation for four scenarios (involving three distinct CGKAs). In each scenario, the leftmost user refreshes their cryptographic keys by broadcasting a commit message. Each figure highlights the nodes for which the commit message contains a public key (○) or a ciphertext (●). When k ciphertexts encrypt the same message, they are regrouped in a same box (□); when this is the case, we may use a mKEM multi-ciphertext instead of k KEM ciphertexts.

Blanking. When a user j is removed from a group in TreeKEM, in order to preserve the tree invariant (TI), all the nodes in their path are blanked out, meaning that these nodes are now considered empty. A blank node becomes populated again if: (i) a user i refreshes their key by broadcasting a commit message and (ii) the user i has the blank node in their path. Note that, due to the tree invariant (TI), a user may not generate a node secret for a node that is not in their co-path. By blanking up to $\lceil \log N \rceil$ nodes, removing users may disrupt the topology of the ratchet tree. This is illustrated by Fig. 5b: the ninth user (from the left) has been removed from the group and their direct path blanked out. When the leftmost user sends a commit message, they now need to send 6 ciphertexts instead of 4, since the node secret of the root needs to be encrypted to 3 nodes instead of 1. In large groups, blank nodes can have a significant adverse effect on the efficiency of TreeKEM.

Our first application of mKEMs is TreeKEM. When implemented with a mKEM instead of a KEM, the efficiency of TreeKEM is more resilient to topology changes provoked by removing users. For the example of Fig. 5b, switching from Kyber to mKyber decreases the (m)KEM-related overhead from 12320 to 6176 bytes, a 50% improvement.

m -ary TreeKEM. A second application of mKEMs can be found by leveraging their flexibility to explore the design space of TreeKEM for better trade-offs. Indeed, instead of a binary tree, one could instantiate TreeKEM with an m -ary tree, for $m > 2$. An example is provided in Fig. 5c for $m = 4$. In a commit message for a full ratchet tree, this reduces the number of public keys to $\lceil \log_m N \rceil$ but increases the number of ciphertexts to $(m - 1) \lceil \log_m N \rceil$, so it’s far from obvious that it would be a good trade-off in general.

Fortunately, for each layer of the ratchet tree, the $m - 1$ ciphertexts in this layer all encrypt the same message. This means these $m - 1$ KEM ciphertexts can be replaced by a single mKEM ciphertext for $m - 1$ recipients, which size may be much shorter in practice. For 16 users, switching from (binary) TreeKEM with Kyber to 4-ary TreeKEM with mKyber decreases the (m)KEM-related bandwidth overhead from 6272 to 3232 bytes. A more detailed presentation of m -ary TreeKEM is found in [KKPP20]

Chained CmpPKE. A final application of mKEMs can be found by taking the m -ary TreeKEM idea to the extreme: by setting $m = N$, we get a flat tree. This is the basis of the Chained CmpPKE protocol, illustrated in Fig. 5d; for this example, the bandwidth overhead related to mKyber is $1472 + (N - 1) \cdot 48 = 2192$ bytes, a 65% gain compared to the standard TreeKEM with Kyber.

Note that the bandwidth overhead is linear in N . Fortunately, we can also exploit the fact that mKyber is efficiently decomposable. The uploaded commit message contains a multi-recipient ciphertext that is a N -uple $(u, (v_i)_i)$, but each recipient only needs to download (u, v_i) to decrypt the message. This allows further savings in the overall protocol. More details can be found in [HKP⁺21].

Table 3: Comparison of the bandwidth costs when using Kyber vs mKyber, for the four scenarii in Fig. 5. For clarity, these numbers ignore (m)KEM-independent values that may be included in the commit message⁴.

Reference	Description	Cost with Kyber			Cost with mKyber			
		ek	ct	Overhead	ek	u	v_i	Overhead
Fig. 5a	TreeKEM (full tree)	4	4	6272	4	4	4	6080
Fig. 5b	TreeKEM (blanked)	4	6	12320	4	4	6	6176
Fig. 5c	4-ary TreeKEM	2	6	6208	2	2	6	3232
Fig. 5d	Chained CmpPKE	1	15	7808	1	1	15	2192

5 Application: Broadcast Scenario

The broadcast scenario, in which one sender transmits the same keying material simultaneously to multiple receivers with authentic static public-keys (discussed within NIST SP 800-56A-rev2) remains an important use-case for the cloud. In particular, this is the case for vendors that network a collection of hosts in a high-availability setting that require the synchronization of cryptographic state/keys across hosts.

A common example of such an application is within a fleet of cloud Hardware Security Modules (HMSs) that provides a key management system to generate, manage and distribute key material for encryption and authentication of user data across cloud computing or large cluster environments. In order for the fleet to synchronize state across all other members, it must first establish a key transport mechanism. To achieve this, SP 800-56A-rev2 relaxes the prohibition against the reuse of an ephemeral Diffie-Hellman key pair in broadcast scenarios, such that the key transport sender can use the same ephemeral key pair when establishing key-wrapping keys with the multiple key-transport receivers.

In such broadcast scenarios, key agreement is often performed by first collecting all static public-keys of members in the group (HSMs) into a public-key repository, which can then

⁴In addition to the (m)KEM-related overhead, commit messages also contain additional data such as a signature, a hash, a constant-size header, etc. See [AHKM22, Figure 8] for an estimate.

be distributed among the fleet so that individual key exchanges can be performed by users as needed. As such, the asymmetric cryptography used in key establishment across large groups can bottleneck fleet performance, particularly in settings in which group members are dynamic and members are changing frequently, or in systems that rotate key establishment keys frequently for forward security requirements. Such issues are exacerbated further when considering the requirement for quantum-resistance. Post-quantum KEMs generally have much larger public keys and ciphertexts than their classical counter parts e.g., 800B for a Kyber512 public key vs. 33B for a (compressed) ECDH P-256 public key.

Using an mKEM instead of a KEM mitigates this bottleneck in multiple ways. First, it drastically decreases bandwidth and computation requirements from the sender, asymptotically 16 times with mKyber instead of Kyber. Second, since the goal is to establish a *single* group key, the sender equipped with a regular KEM has to use the KEM-DEM paradigm to wrap the group key. On the other hand, an mKEM already produces a single key. This means that to send a key to N members with Kyber, the sender generates N Kyber ciphertexts and N symmetric encryptions, while with mKyber it generates one mKyber ciphertext and no symmetric encryptions. In other words, we take advantage of the fact that Kyber already consists of an underlying *encryption* scheme where we can choose the key to be the same. This reduces sender and receiver computation and bandwidth cost.

6 Further reading

Due to space constraints, this paper only gives a high-level view of mKEMs and their applications. We provide further references here.

Classical mKEMs. The possibility of saving bandwidth when encrypting to multiple recipients has first been studied in [BBM00, Kur02, BBS03], via multi-recipient constructions based on El Gamal and Cramer-Shoup, which rely on classical assumptions. The term mKEM has been coined by Smart [Sma05] and studied further in works [HK07, BF07, MH13].

Post-quantum mKEMs. Comparatively, the study of post-quantum mKEMs is much more recent. A LPN-based construction was proposed in [CLQY18] but later proven insecure due in [KKPP20] to a misuse of the Fujisaki-Okamoto transform. The first secure constructions for post-quantum mKEMs have been proposed in [KKPP20] and further developed in [HKP⁺21, AHK⁺23]. These three papers propose multi-recipient adaptations of the LPR/Lindner-Peikert framework [LPR10, LP11], which underlied many of the candidates to the 2017 NIST PQC call: Kyber (future ML-KEM), FrodoKEM, Saber, etc.⁵

Applications. The works of [KKPP20, HKP⁺21, AHKM22] apply mKEMs in the context of secure messaging, particularly the MLS protocol [BBR⁺23]. The present work proposes a natural application in the context of broadcast scenarios (Section 5), and we expect further applications to be found.

References

- [ABH⁺21] Joël Alwen, Bruno Blanchet, Eduard Hauck, Eike Kiltz, Benjamin Lipp, and Doreen Riepel. Analysing the HPKE standard. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 87–116. Springer, Heidelberg, October 2021.
- [AHK⁺23] Joël Alwen, Dominik Hartmann, Eike Kiltz, Marta Mularczyk, and Peter Schwabe. Post-quantum multi-recipient public key encryption. In Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda, editors, *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS 2023, Copenhagen, Denmark, November 26-30, 2023*, pages 1108–1122. ACM, 2023.

⁵[KKPP20] also proposed mKEMs based on SIDH/SIKE and CSIDH; the recent cryptanalytic advances on CSIDH [BS20, Pei20] and SIDH [CD23, MMP⁺23, Rob23] make these insecure.

- [AHKM22] Joël Alwen, Dominik Hartmann, Eike Kiltz, and Marta Mularczyk. Server-aided continuous group key agreement. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 69–82. ACM Press, November 2022.
- [BBKS07] Mihir Bellare, Alexandra Boldyreva, Kaoru Kurosawa, and Jessica Staddon. Multirecipient encryption schemes: How to save on bandwidth and computation without sacrificing security. *IEEE Transactions on Information Theory*, 53(11):3927–3943, 2007.
- [BBLW22] Richard Barnes, Karthikeyan Bhargavan, Benjamin Lipp, and Christopher A. Wood. Hybrid Public Key Encryption. RFC 9180, February 2022.
- [BBM00] Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 259–274. Springer, Heidelberg, May 2000.
- [BBR⁺23] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. The Messaging Layer Security (MLS) Protocol. RFC 9420, July 2023.
- [BBS03] Mihir Bellare, Alexandra Boldyreva, and Jessica Staddon. Randomness re-use in multi-recipient encryption schemes. In Yvo Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 85–99. Springer, Heidelberg, January 2003.
- [BF07] Manuel Barbosa and Pooya Farshim. Randomness reuse: Extensions and improvements. In *IMA International Conference on Cryptography and Coding*, pages 257–276. Springer, 2007.
- [BS20] Xavier Bonnetain and André Schrottenloher. Quantum security analysis of CSIDH. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 493–522. Springer, Heidelberg, May 2020.
- [CD23] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 423–447. Springer, Heidelberg, April 2023.
- [CLQY18] Haitao Cheng, Xiangxue Li, Haifeng Qian, and Di Yan. Cca secure multi-recipient kem from lpn. In *ICICS*, pages 513–529. Springer, 2018.
- [FO13] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *Journal of Cryptology*, 26(1):80–101, January 2013.
- [GW09] Craig Gentry and Brent Waters. Adaptive security in broadcast encryption systems (with short ciphertexts). In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 171–188. Springer, Heidelberg, April 2009.
- [HBD⁺22] Andreas Hülsing, Daniel J. Bernstein, Christoph Dobraunig, Maria Eichlseder, Scott Fluhrer, Stefan-Lukas Gazdag, Panos Kampanakis, Stefan Kölbl, Tanja Lange, Martin M. Lauridsen, Florian Mendel, Ruben Niederhagen, Christian Rechberger, Joost Rijneveld, Peter Schwabe, Jean-Philippe Aumasson, Bas Westerbaan, and Ward Beullens. SPHINCS⁺. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [HHK17] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 341–371. Springer, Heidelberg, November 2017.
- [HK07] Dennis Hofheinz and Eike Kiltz. Secure hybrid encryption from weakened key encapsulation. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 553–571. Springer, Heidelberg, August 2007.
- [HKP⁺21] Keitaro Hashimoto, Shuichi Katsumata, Eamonn Postlethwaite, Thomas Prest, and Bas Westerbaan. A concrete treatment of efficient continuous group key agreement via multi-recipient PKEs. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021*, pages 1441–1462. ACM Press, November 2021.

-
- [KKPP20] Shuichi Katsumata, Kris Kwiatkowski, Federico Pintore, and Thomas Prest. Scalable ciphertext compression techniques for post-quantum KEMs and their applications. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part I*, volume 12491 of *LNCS*, pages 289–320. Springer, Heidelberg, December 2020.
- [Kur02] Kaoru Kurosawa. Multi-recipient public-key encryption with shortened ciphertext. In David Naccache and Pascal Paillier, editors, *PKC 2002*, volume 2274 of *LNCS*, pages 48–63. Springer, Heidelberg, February 2002.
- [LDK⁺22] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In Aggelos Kiayias, editor, *CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339. Springer, Heidelberg, February 2011.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Heidelberg, May / June 2010.
- [MH13] Takahiro Matsuda and Goichiro Hanaoka. Key encapsulation mechanisms from extractable hash proof systems, revisited. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 332–351. Springer, Heidelberg, February / March 2013.
- [MMP⁺23] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 448–471. Springer, Heidelberg, April 2023.
- [Pei20] Chris Peikert. He gives C-sieves on the CSIDH. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 463–492. Springer, Heidelberg, May 2020.
- [PFH⁺22] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [PPS14] Alexandre Pinto, Bertram Poettering, and Jacob C. N. Schuldt. Multi-recipient encryption, revisited. In Shiho Moriai, Trent Jaeger, and Kouichi Sakurai, editors, *ASIACCS 14*, pages 229–238. ACM Press, June 2014.
- [Rob23] Damien Robert. Breaking SIDH in polynomial time. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 472–503. Springer, Heidelberg, April 2023.
- [SAB⁺22] Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [Sma05] Nigel P. Smart. Efficient key encapsulation to multiple parties. In Carlo Blundo and Stelvio Cimato, editors, *SCN 04*, volume 3352 of *LNCS*, pages 208–219. Springer, Heidelberg, September 2005.