Covercrypt: an Efficient Early-Abort KEM
for Hidden Access Policies with Traceability
from the DDH and LWE

Théophile Brézot[1], Paola de Perthuis[1,2], and David Pointcheval[2]

[1] Cosmian, Paris, France
[2] DIENS, École normale supérieure, université PSL, CNRS, Inria, Paris, France

Abstract. Attribute-Based Encryption (ABE) is a very attractive primitive to limit access according to specific rights. While very powerful instantiations have been offered, under various computational assumptions, they rely on either classical or post-quantum problems, and are quite intricate to implement, generally resulting in poor efficiency; the construction we offer results in a powerful efficiency gap with respect to existing solutions. With the threat of quantum computers, post-quantum solutions are important, but not yet tested enough to rely on such problems only. We thus first study an hybrid approach to rely on the best of the two worlds: the scheme is secure if at least one of the two underlying assumptions is still valid (i.e. the DDH and LWE). Then, we address the ABE problem, with a practical solution delivering encrypted contents such that only authorized users can decrypt, without revealing the target sets, while also granting tracing capabilities. Our scheme is inspired by the Subset Cover framework where the users' rights are organized as subsets and a content is encrypted with respect to a subset covering of the target set. Quite conveniently, we offer black-box modularity: one can easily use any public-key encryption of their choice, such as Kyber, with their favorite library, to combine it with a simple ElGamal variant of key encapsulation mechanisms, providing strong security guarantees.

1 Introduction

Key Encapsulation Mechanisms (KEM) enable the transmission of symmetric keys at the beginning of an interaction while retaining trust that only the intended recipient will be able to get access to this encapsulated key. Once this trusted transmission has been established, users can privately communicate using this encapsulated secret key with the advantages of symmetric encryption, granting compact ciphertexts of similar size as corresponding cleartexts. Namely, they can be used to build Public-Key Encryption (PKE) schemes in the KEM-DEM (for Data Encapsulation Mechanism) paradigm [Sho01].

In organizations with complex structures, one will want to have more functionalities, namely being able to share a key among all users verifying a policy on a set of attributes, all at once. To this aim, KEMs constructed out of Attribute-Based Encryption (ABE) have been designed, in which keys can be encapsulated by being encrypted with these schemes for which all users verifying the specified attributes policy will be able to decrypt and thus decapsulate the key. These ABE primitives (stemming from [GPSW06]) are very powerful as they can cover any possible logical combination of the attributes, however this comes at an efficiency cost, and for practical use-cases, one will only need to encrypt for some of these existing combinations, for a limited number of attributes; this work is in this setting's scope, in which one can actually replace ABE constructions with encryption with respect to a union of attribute subsets. In these use-cases, it can also be relevant to get anonymity, meaning that a user should never know for which policy a ciphertext was produced, except if it is the policy they are using to successfully decrypt. In the case of ABE, this is called *attribute hiding*. This can also be used to get anonymous authentication (for instance in mobile network contexts) to service providers sending encapsulations without users needing to send out requests that would identify them.

Additionally, with current preoccupations with respect to the threat of quantum computers on classical cryptography, granting resistance to these for data that needs to be kept private on the long term is becoming a necessity. However, post-quantum cryptographic schemes are newer and only beginning to be used, one should try to keep current schemes' security properties. In fact, several security agencies are handing out guidelines for pre- and post-quantum security hybridization, meaning that cryptographic schemes should retain all their security properties even if one of the two pre- or post-quantum schemes is broken.

Another area of interest in this context in which users share some common keys, is the ability to still identify them uniquely, in case they choose to send some of there decapsulation capabilities to another

party. Thus, if someone leaks some secret information they were supposed to keep to themselves, we would like to trace these so-called traitors, with *traceability*.

Related Work. This work combines many desirable properties for the use of KEMs in practical contexts, that other previous works had not, and since it covers only the practical contexts in which one would wish for ABE-based constructions, it compares favorably in efficiency with respect to such post-quantum schemes built from ABE, in addition with providing traceability and post- and pre-quantum hybridization.

Anonymous Broadcast Encryption. Our simplified access structure with strong privacy has a similar flavor as previous works [LPQ12, FP12, LG18] on broadcast encryption with anonymity, with optimizations on the decryption time. However, they do not handle black-box post-quantum security nor traceability.

Post-Quantum Key-Policy ABE. Then, providing post-quantum resistance, the closest related works are Key-Policy ABEs (KP-ABE) based on LWE. Some theoretical works such as [Wee21] provide results with good asymptotic bounds, but are unsuited for use with practical parameters, and others, like [DDP+17], provide implementable results, but even with their comparable lowest policy circuit depth, their encryption time is about a hundred times bigger than ours, their decryption time about ten times bigger, and their RLWE parameters lead to bigger ciphertext sizes than ours. Also, they do not provide anonymity nor traceability.

Hybridization for Pre and Post-Quantum Security. Our work, in the line of security agency and standardization organizations recommendations, enables the hybridization of both pre- and post-quantum schemes, so that its security holds if either one of the schemes do. The use of the post-quantum scheme is totally black-box, enabling combinations with other semantically secure public-key encryption schemes. This is in the line of previous work to combine KEMs to get the best security out of the individual ones combined, such as [GHP18], and in [BBF+19], where the specific problem of combining pre- and post-quantum schemes against various types of classical or quantum adversaries was studied.

Our Contributions. Our final instantiation called Covercrypt provides an efficient KEM for hidden access policies with traceability, ensuring both pre- and post-quantum securities, along with a Rust implementation of the scheme[3].

An Efficient KEM with Hidden Access Policies. Our scheme provides efficiency with respect to the state-of-the-art in post-quantum KP-ABE schemes by restricting its scope to depth-one policy circuits. The attributes for which a key is encapsulated are kept hidden, providing anonymity. Also, we gain time on the decryption with an early-abort paradigm, in which one can quickly test whether a ciphertext was encrypted for one of their attributes, using a tag, and retaining the anonymity properties of the scheme. Our ciphertexts are of size $96 + \#B \times 1088$ Bytes, where B is the list of attribute-subsets the key is encapsulated for. On the other hand, user's keys are of size $(\#A + 1) \times 64$ Bytes, where A is the list of attributes for the user. For $\#B$ ranging from 1 to 5, encapsulation takes from 350 to 950 microseconds, and decapsulation, from 230 to 480 microseconds, with an affine dependency in the user's attributes (see Section 7).

Traceability. As an optional feature, the pre-quantum ElGamal part of our scheme provides traceability under the Decisional Diffie-Hellman (DDH) assumption. It makes sense to consider traceability with pre-quantum security as this is a short-term security requirement, if users are currently misbehaving, whereas the post-quantum security preserves the privacy property, which is important on the long-term, as ciphertexts can be stored until their security is broken in the future. Our implementation covers the case were traitors do not collude; we also show how the scheme can be instantiated for arbitrarily t-large collusions, but the tracing time then grows exponentially in t. A KEM can be used to broadcast symmetric encryption keys, but also for authentication, and in such an interactive context, implementing tracing requests is easily done in practice.

---

[3] `https://github.com/Cosmian/cover_crypt`

2 Definitions

Public-Key Encryption (PKE) allows the transmission of hidden information that only the intended recipient will be able to uncover. To make the scheme independent of the format of the cleartext message, the usual paradigm for encryption is the KEM-DEM [Sho01], where one first encapsulates a session key that only the recipient can recover, and then encrypts the payload under that key. The former step uses a Key Encapsulation Mechanism (KEM) and the latter a Data Encapsulation Mechanism (DEM), that is usually instantiated with an Authenticated Encryption, such as AES256-GCM[4], that provides both privacy and authenticity of plaintexts. As our work is built from KEMs, we hereafter recall some formal definitions.

Notations. Henceforth, many security notions will be characterized by the computational indistinguishability between two distributions $\mathcal{D}_0$ and $\mathcal{D}_1$. It will be measured by the advantage an adversary $\mathcal{A}$ can have in distinguishing them:

$$\mathsf{Adv}(\mathcal{A}) = \Pr_{\mathcal{D}_1}[\mathcal{A}(x) = 1] - \Pr_{\mathcal{D}_0}[\mathcal{A}(x) = 1] = 2 \times \Pr_{\mathcal{D}_b}[\mathcal{A}(x) = b] - 1.$$

Then, we will denote $\mathsf{Adv}(\tau)$ the maximal advantage over all the adversaries with running-time bounded by τ. A first pair of distributions is used in the famous ElGamal encryption scheme, with Diffie-Hellman tuples in $\mathbb{G} = \langle g \rangle$, a group of prime order p, spanned by a generator g, and denoted multiplicatively:

Definition 1 (Decisional Diffie-Hellman Problem). *The DDH assumption in a group $\mathbb{G}$ (DDH$_\mathbb{G}$) of prime order p, with a generator g, states that the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ are computationally hard to distinguish, where*

$$\mathcal{D}_0 = \{(g^a, g^b, g^{ab}), a, b \xleftarrow{\$} \mathbb{Z}_p\} \qquad\qquad \mathcal{D}_1 = \{(g^a, g^b, g^c), a, b, c \xleftarrow{\$} \mathbb{Z}_p\}$$

and we will denote $\mathsf{Adv}^{\mathsf{ddh}}_\mathbb{G}(\mathcal{A})$ the advantage of an adversary $\mathcal{A}$.

When studying the Kyber post-quantum encryption scheme, we will also need another algebraic structure, with indistinguishable distributions. We will denote $\mathsf{R} = \mathbb{Z}[X]/(X^n + 1)$ (resp. $\mathsf{R}_q = \mathbb{Z}_q[X]/(X^n + 1)$) the ring of polynomials of degree at most $n-1$ with integer coefficients (resp. with coefficients in $\mathbb{Z}_q$, for a small prime q). We take n as power of 2, where $X^n + 1$ is the $\frac{n}{2}$-th cyclotomic polynomial. We denote $\mathcal{B}_\eta$ the centered binomial distribution of parameter η. When a polynomial is sampled according to $\mathcal{B}_\eta$, it means each of its coefficient is sampled from that distribution. We will also use vectors $\mathbf{e} \in \mathsf{R}_q^k$ and matrices $\mathbf{A} \in \mathsf{R}_q^{m \times k}$ in R_q:

Definition 2 (Decisional Module Learning-with-Error Problem). *The DMLWE assumption in R_q (DMLWE$_{\mathsf{R}_q, m, k, \eta}$) states that the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ are computationally hard to distinguish, where*

$$\mathcal{D}_0 = \{(\mathbf{A}, \mathbf{b}), \mathbf{A} \xleftarrow{\$} \mathsf{R}_q^{m \times k}, (\mathbf{s}, \mathbf{e}) \xleftarrow{\$} \mathcal{B}_\eta^k \times \mathcal{B}_\eta^m, \mathbf{b} \leftarrow \mathbf{As} + \mathbf{e}\}$$
$$\mathcal{D}_1 = \{(\mathbf{A}, \mathbf{b}), \mathbf{A} \xleftarrow{\$} \mathsf{R}_q^{m \times k}, \mathbf{b} \xleftarrow{\$} \mathcal{B}_\eta^m\}$$

We will denote $\mathsf{Adv}^{\mathsf{dmlwe}}_{\mathsf{R}_q, m, k, \eta}(\mathcal{A})$ the advantage of an adversary $\mathcal{A}$.

Pseudorandom Generators (PRG). A long line of cryptographic works consider PRGs [HILL99, App12], as one of the theoretical foundations of modern cryptography. A PRG $\mathsf{PRG} : \{0; 1\}^\mu \to \{0; 1\}^\nu$ is deterministic function which should have the property that uniformly distributed inputs on $\{0; 1\}^\mu$ should have outputs through PRG indistinguishable from uniformly random samples of $\{0; 1\}^\nu$ with respect to a PPT adversary. The bigger ν is with respect to μ, the more challenging constructing such a PRG becomes. We define a PRG's security as:

Definition 3 (IND-security of a PRG). *Let $\mathsf{PRG} : \{0; 1\}^\mu \to \{0; 1\}^\nu$ be a deterministic function. Then PRG is an IND-secure PRG if the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ are computationally hard to distinguish, where*

$$\mathcal{D}_0 = \{y, x \xleftarrow{\$} \{0; 1\}^\mu, y \leftarrow \mathsf{PRG}(x)\} \qquad\qquad \mathcal{D}_1 = \{y, y \xleftarrow{\$} \{0; 1\}^\nu\}$$

We will denote $\mathsf{Adv}^{\mathsf{ind}}_{\mathsf{PRG}_{\mu, \nu}}(\mathcal{A})$ the advantage of an adversary $\mathcal{A}$.

---

[4] https://docs.rs/aes-gcm/latest/aes_gcm/

Key Encapsulation Mechanism. A Key Encapsulation Mechanism KEM is defined by three algorithms:

- KEM.KeyGen(1^κ): the *key generation algorithm* outputs a pair of public and secret keys $(\mathsf{pk}, \mathsf{sk})$;
- KEM.Enc(pk): the *encapsulation algorithm* generates a session key K and an encapsulation C of it, and outputs the pair (C, K);
- KEM.Dec(sk, C): the *decapsulation algorithm* outputs the key K encapsulated in C.

Correctness. A correct KEM satisfies $\mathsf{Adv}^{\mathsf{cor}}_{\mathsf{KEM}}(\kappa) = 1 - \Pr_{\mathcal{D}}[\mathsf{Ev}] = \mathsf{negl}(\kappa)$, for

$$\mathcal{D} = \{(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KEM.KeyGen}(1^\kappa), (C, K) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk}) : (\mathsf{sk}, C, K)\}$$
$$\mathsf{Ev} = [\mathsf{KEM.Dec}(\mathsf{sk}, C) = K]$$

Session-Key Privacy. On the other hand, such a KEM is said to provide *session-key privacy* (denoted SK-IND) in the key space $\mathcal{K}$, if the encapsulated key is indistinguishable from a random key in $\mathcal{K}$. More formally, a KEM is SK-IND-secure if for any adversary $\mathcal{A}$, $\mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}) = \mathsf{negl}(\kappa)$, in distinguishing $\mathcal{D}_0$ and $\mathcal{D}_1$, where

$$\mathcal{D}_b = \left\{ \begin{array}{l} (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KEM.KeyGen}(1^\kappa), \\ (C, K_0) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk}), K_1 \xleftarrow{\$} \mathcal{K} \end{array} : (\mathsf{pk}, C, K_b) \right\}$$

Public-Key Privacy. One can additionally expect anonymity of the receiver, also known as *public-key privacy* (denoted PK-IND), if the encapsulation does not leak any information about the public key, first defined in [BBDP01]. More formally, a KEM is PK-IND-secure if for any adversary $\mathcal{A}$, $\mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}) = \mathsf{negl}(\kappa)$, in distinguishing $\mathcal{D}_0$ and $\mathcal{D}_1$, where

$$\mathcal{D}_b = \left\{ \begin{array}{l} \text{For } i = 0, 1: \\ \quad (\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KEM.KeyGen}(1^\kappa), \\ \quad (C_i, K_i) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk}_i) \end{array} : (\mathsf{pk}_0, \mathsf{pk}_1, C_b) \right\}$$

ElGamal-based KEM. In a group $\mathbb{G}$ of prime order p, with a generator g:

- EG.KeyGen(1^κ): sample random $\mathsf{sk} = x \xleftarrow{\$} \mathbb{Z}_p$ and set $\mathsf{pk} = h \leftarrow g^x$;
- EG.Enc(pk): sample a random $r \xleftarrow{\$} \mathbb{Z}_p$ and set $C \leftarrow g^r$ together with $K \leftarrow h^r$;
- EG.Dec(sk, C): output $K \leftarrow C^x$.

Under the DDH assumption in $\mathbb{G}$, this KEM is both SK-IND and PK-IND with $\mathcal{K} = \mathbb{G}$. The formal security proofs for an extended version of this scheme will be given later, we thus postpone the analysis of this scheme.

Key Encapsulation Mechanism with Access Control. A KEM with Access Control allows multiple users to access the encapsulated key K from C, according to a rule $\mathcal{R}$ applied on X in the user's key usk and Y in the encapsulation C. It is defined by four algorithms:

- KEMAC.Setup(1^κ) outputs the global public parameters PK and the master secret key MSK;
- KEMAC.KeyGen(MSK, Y) outputs the user's secret key usk according to Y;
- KEMAC.Enc(PK, X) generates a session key K and an encapsulation C of it according to X;
- KEMAC.Dec(usk, C) outputs the key K encapsulated in C.

Correctness. A KEMAC is correct if $\mathsf{Adv}^{\mathsf{cor}}_{\mathsf{KEMAC}}(\kappa) = 1 - \Pr_{\mathcal{D}}[\mathsf{Ev}] = \mathsf{negl}(\kappa)$, for

$$\mathcal{D} = \left\{ \begin{array}{l} \forall (X, Y) \text{ such that } \mathcal{R}(X, Y) = 1, \\ (\mathsf{PK}, \mathsf{MSK}) \leftarrow \mathsf{KEMAC.KeyGen}(1^\kappa), \\ \mathsf{usk} \leftarrow \mathsf{KEMAC.KeyGen}(\mathsf{MSK}, Y), \\ (C, K) \leftarrow \mathsf{KEMAC.Enc}(\mathsf{PK}, X) \end{array} : (\mathsf{usk}, C, K) \right\}$$
$$\mathsf{Ev} = [\mathsf{KEMAC.Dec}(\mathsf{usk}, C) = K].$$

Session-Key Privacy. As for the basic KEM, one may expect some privacy properties. Session-key privacy is modeled by indistinguishability of ciphertexts, even if the adversary has received some decryption keys, as soon as associated Y_i are incompatible with X ($\mathcal{R}(X, Y_i) = 0$). Such a KEMAC is said to be SK-IND-secure in the key space $\mathcal{K}$ if for any adversary $\mathcal{A}$, that can ask any key usk_i, using oracle $\mathcal{O}\mathsf{KeyGen}(Y_i)$ that stores Y_i in the set $\mathcal{Y}$ and outputs $\mathsf{KEMAC.KeyGen}(\mathsf{MSK}, Y_i)$, $\mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEMAC}}(\mathcal{A}) = \mathsf{negl}(\kappa)$, for $b \overset{\$}{\leftarrow} \{0; 1\}$ and

$$\mathcal{D}_b = \left\{ \begin{array}{l} (\mathsf{PK}, \mathsf{MSK}) \leftarrow \mathsf{KEMAC.Setup}(1^\kappa), \\ (\mathsf{state}, X) \leftarrow \mathcal{A}^{\mathcal{O}\mathsf{KeyGen}(\cdot)}(\mathsf{PK}), \\ (C, K_0) \leftarrow \mathsf{KEMAC.Enc}(\mathsf{PK}, X), K_1 \overset{\$}{\leftarrow} \mathcal{K} \end{array} : (\mathsf{state}, C, K_b) \right\}$$

$$\mathsf{BadXY} = [\exists Y_i \in \mathcal{Y}, \mathcal{R}(X, Y_i) = 1]$$

$$\mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEMAC}}(\mathcal{A}) = 2 \times \Pr_{\mathcal{D}_b}[\mathcal{A}^{\mathcal{O}\mathsf{KeyGen}(\cdot)}(\mathsf{state}, C, K_b) = b \,|\, \neg\mathsf{BadXY}] - 1.$$

We note the bad event BadXY (decided at the end of the game) should be avoided by the adversary, as it reduces its advantage: this indeed leads to a trivial guess, and this is considered as a non-legitimate attack.

Access-Control Privacy. In addition, one could want to hide the parameter X used in the encapsulation C even if the adversary $\mathcal{A}$ can ask any key usk_i for Y_i such that $\mathcal{R}(X_0, Y_i) = \mathcal{R}(X_1, Y_i) = 0$ for all i, using oracle $\mathcal{O}\mathsf{KeyGen}(Y_i)$ that stores Y_i in the set $\mathcal{Y}$ and outputs $\mathsf{KEMAC.KeyGen}(\mathsf{MSK}, Y_i)$. A KEMAC is said to be AC-IND-secure if for any adversary $\mathcal{A}$, that can ask any key usk_i, using oracle $\mathcal{O}\mathsf{KeyGen}(Y_i)$ that stores Y_i in the set $\mathcal{Y}$ and outputs $\mathsf{KEMAC.KeyGen}(\mathsf{MSK}, Y_i)$, $\mathsf{Adv}^{\mathsf{ac\text{-}ind}}_{\mathsf{KEMAC}}(\mathcal{A}) = \mathsf{negl}(\kappa)$, for $b \overset{\$}{\leftarrow} \{0; 1\}$ and

$$\mathcal{D}_b = \left\{ \begin{array}{l} (\mathsf{PK}, \mathsf{MSK}) \leftarrow \mathsf{KEMAC.Setup}(1^\kappa), \\ (\mathsf{state}, X_0, X_1) \leftarrow \mathcal{A}^{\mathcal{O}\mathsf{KeyGen}(\cdot)}(\mathsf{PK}), \\ (C_i, K_i) \leftarrow \mathsf{KEMAC.Enc}(\mathsf{PK}, X_i), \text{ for } i = 0, 1 \end{array} : (\mathsf{state}, C_b) \right\}$$

$$\mathsf{BadXY} = [\exists Y_i \in \mathcal{Y}, \mathcal{R}(X_0, Y_i) = 1 \vee \mathcal{R}(X_1, Y_i) = 1]$$

$$\mathsf{Adv}^{\mathsf{ac\text{-}ind}}_{\mathsf{KEMAC}}(\mathcal{A}) = 2 \times \Pr_{\mathcal{D}_b}[\mathcal{A}^{\mathcal{O}\mathsf{KeyGen}(\cdot)}(\mathsf{state}, C_b) = b \,|\, \neg\mathsf{BadXY}] - 1,$$

where we again condition the advantage to legitimate attacks only.

Traceability. In any multi-user setting, to avoid abuse of the decryption keys, one may want to be able to trace a user (or their personal key) from the decryption mechanism, and more generally from any *useful* decoder, either given access to the key material in the device (white-box tracing) or just interacting with the device (black-box tracing). Without any keys, one expects session-key privacy, but as soon as one knows a key, one can distinguish the session-key. Then, we will call a *useful* pirate decoder $\mathcal{P}$ a good distinguisher against session-key privacy, that behaves differently with the real and a random key. But of course, this pirate decoder can be built from multiple user' keys, called traitors, and one would like to be able to trace at least one of them.

A weaker variant of traceability is just a confirmation of candidate traitors, and we will target this goal: if a pirate decoder $\mathcal{P}$ has been generated from a list $\mathcal{T} = \{Y_i\}$ of traitors' keys, a confirmer algorithm $\mathcal{C}$ can output, from a valid guess $\mathcal{G}$ for $\mathcal{T}$, at least one traitor in $\mathcal{T}$. More formally, let us consider any adversary $\mathcal{A}$ that can ask for key generation through oracle $\mathcal{O}\mathsf{KeyGen}(Y_i)$, that gets $\mathsf{usk}_i \leftarrow \mathsf{KEMAC.KeyGen}(\mathsf{MSK}, Y_i)$, outputs nothing but appends the new user Y_i in $\mathcal{U}$, and then corrupt some users through the corruption oracle $\mathcal{O}\mathsf{Corrupt}(Y_i)$, that outputs usk_i and appends Y_i in $\mathcal{T}$, to build a *useful* pirate decoder $\mathcal{P}$, then there is a *correct* confirmer algorithm $\mathcal{C}$ that outputs a traitor T, with *negligible error*: for $b \overset{\$}{\leftarrow} \{0; 1\}$ and

$$\mathcal{D} = \left\{ \begin{array}{l} (\mathsf{PK}, \mathsf{MSK}) \leftarrow \mathsf{KEMAC.Setup}(1^\kappa), \mathcal{P} \leftarrow \mathcal{A}^{\mathcal{O}\mathsf{KeyGen}(\cdot), \mathcal{O}\mathsf{Corrupt}(\cdot)}(\mathsf{PK}), \\ X \text{ such that } \forall Y_i \in \mathcal{T}, \mathcal{R}(X, Y_i) = 1, \\ (C, K_0) \leftarrow \mathsf{KEMAC.Enc}(\mathsf{PK}, X), K_1 \overset{\$}{\leftarrow} \mathcal{K} : \\ \qquad\qquad\qquad (\mathsf{MSK}, \mathcal{P}, \mathcal{U}, \mathcal{T}, C, K_0, K_1) \end{array} \right\},$$

we denote

- $\mathcal{P}$ to be useful, if $2 \times \Pr_{\mathcal{D}, b}[\mathcal{P}(C, K_b) = b] - 1$ is non-negligible;
- $\mathcal{C}$ to be correct, if $\Pr_{\mathcal{D}}[T \in \mathcal{T} \,|\, T \leftarrow \mathcal{C}^{\mathcal{P}(\cdot, \cdot)}(\mathsf{MSK}, \mathcal{T})]$ is overwhelming;
- $\mathcal{C}$ is error-free if for any $\mathcal{G} \subset \mathcal{U}$, $\Pr_{\mathcal{D}}[T \notin \mathcal{T} \,|\, T \leftarrow \mathcal{C}^{\mathcal{P}(\cdot, \cdot)}(\mathsf{MSK}, \mathcal{G}) \wedge T \neq \bot]$ is negligible.

More concretely, we say that the decoder $\mathcal{P}$ is *useful* if it can distinguish the real key from a random key with significant advantage. Then, from such a useful decoder, the confirmer $\mathcal{C}$ is *correct* if it outputs a traitor with overwhelming probability, when it starts from the correct set $\mathcal{T}$ of candidates. Eventually, it should be *error-free*: $\mathcal{T}$ does not output an honest user, but with negligible probability. The t-confirmation limits the number of corrupted users in $\mathcal{T}$ to t.

Hybrid KEM. While one can never exclude an attack against a cryptographic scheme, combining several independent approaches reduces the risks. This is the way one suggests to apply post-quantum schemes, in combination with classical schemes, in order to be sure to get the best security.

Hybrid KEM Construction. Let us first study the combination of two KEMs (KEM_1 and KEM_2), so that as soon as one of them achieves SK-IND security, the hybrid KEM achieves SK-IND security too.

We need both KEMs to generate keys in $\mathcal{K}$, with a group structure and internal law denoted $\oplus$. One can also find it in appendix H, figure 2:

- $\mathsf{KEM.KeyGen}(1^\kappa)$ calls $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KEM}_i.\mathsf{KeyGen}(1^\kappa)$, for $i \in \{1,2\}$ and outputs $\mathsf{pk} \leftarrow (\mathsf{pk}_1, \mathsf{pk}_2)$ and $\mathsf{sk} \leftarrow (\mathsf{sk}_1, \mathsf{sk}_2)$;
- $\mathsf{KEM.Enc}(\mathsf{pk})$ parses pk as $(\mathsf{pk}_1, \mathsf{pk}_2)$, calls $(C_i, K_i) \leftarrow \mathsf{KEM}_i.\mathsf{Enc}(\mathsf{pk}_i)$ for $i \in \{1,2\}$, and outputs $(C = (C_1, C_2), K = K_1 \oplus K_2)$;
- $\mathsf{KEM.Dec}(\mathsf{sk}, C)$ parses sk as $(\mathsf{sk}_1, \mathsf{sk}_2)$ and C as (C_1, C_2), then calls both $K_i \leftarrow \mathsf{KEM}_i.\mathsf{Dec}(\mathsf{sk}_i, C_i)$, and outputs $K = K_1 \oplus K_2$.

Security Properties. As expected, we can prove that as soon as one of them achieves SK-IND security, the hybrid KEM achieves SK-IND security too. This also follows from [GHP18]'s first lemma. However, for PK-IND security of KEM, we need both the underlying schemes to be PK-IND secure. This second property is not as crucial as the first one: none of the other security properties we show for the schemes depend on it, and here the only property at stake is the anonymity of the receiver of the encapsulated keys, not the keys themselves. The proofs of theorems 4 and 5 can be found in appendices D.4 and D.5, respectively.

Theorem 4 (Session-Key Privacy). *If at least one of the underlying KEM_1 and KEM_2 is SK-IND-secure, the hybrid KEM is SK-IND-secure:*

$$\mathsf{Adv}_{KEM}^{sk\text{-}ind}(\tau) \leq \min\{\mathsf{Adv}_{KEM_1}^{sk\text{-}ind}(\tau), \mathsf{Adv}_{KEM_2}^{sk\text{-}ind}(\tau)\}.$$

Theorem 5 (Public-Key Privacy). *If both underlying KEM_1 and KEM_2 are PK-IND-secure, the hybrid KEM is PK-IND-secure:*

$$\mathsf{Adv}_{KEM}^{pk\text{-}ind}(\tau) \leq \mathsf{Adv}_{KEM_1}^{pk\text{-}ind}(\tau) + \mathsf{Adv}_{KEM_2}^{pk\text{-}ind}(\tau).$$

We will also use Public-Key Encryption (PKE), which is recalled in the Appendix A.

3 Authenticated Key Encapsulation Mechanism

With public-key privacy, one cannot know who is the actual receiver, and needs to check the decapsulated session key with an authenticated encryption scheme to know whether they were a recipient or not. The latter check can be time-consuming when applied on a large data content (or when there are multiple decryption keys to try). We can hope to have quick key confirmation, if the additional *Authentication* (AUTH) property is satisfied.

Authentication. A KEM provides *authentication* (denoted AUTH) if it satisfies $\mathsf{Adv}_{KEM}^{auth}(\kappa) = 1 - \Pr_{\mathcal{D}}[\mathsf{Ev}] = \mathsf{negl}(\kappa)$, for

$$\mathcal{D} = \left\{ \begin{array}{l} \forall i \in \{0;1\}, (\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KEM.KeyGen}(1^\kappa), \\ (C, K) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk}_0) : (\mathsf{sk}_1, C) \end{array} \right\}$$

$$\mathsf{Ev} = [\mathsf{KEM.Dec}(\mathsf{sk}_1, C) = \bot].$$

We stress this is a weak authentication definition, but strong enough for our further early-abort technique. We indeed just want to exclude a ciphertext to be valid under two keys, at random. There is no malicious behavior.

We present a generic conversion to add the AUTH property to any KEM, while retaining previous properties (SK-IND and PK-IND). To this aim, we use a hash function $\mathcal{H}$. In the security analysis, it will be modeled by a random oracle that outputs a new random and independent bitstring for any new query.

Key Encapsulation Mechanisms with Authentication. We present below (and in the Appendix H in Figure 1) a KEM' with authentication from a KEM that outputs κ-bit keys, with two security parameters: k, the length of the new encapsulated key, and ℓ, the length of the verification tag. We also use a PRG $\mathsf{PRG} : \{0;1\}^\kappa \to \{0;1\}^{k+\ell}$. We require that in $\mathsf{KEM}.\mathsf{Enc}$'s outputs (C,K), with K looking uniform in $\{0;1\}^\kappa$.

- $\mathsf{KEM}'.\mathsf{KeyGen}(1^\kappa)$ runs $(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{KEM}.\mathsf{KeyGen}(1^\kappa)$;
- $\mathsf{KEM}'.\mathsf{Enc}(\mathsf{pk})$ runs $(c,s) \leftarrow \mathsf{KEM}.\mathsf{Enc}(\mathsf{pk})$ and gets $U\|V \leftarrow \mathsf{PRG}(s)$. One then outputs $C \leftarrow (c,V)$ together with the encapsulated key $K \leftarrow U$;
- $\mathsf{KEM}'.\mathsf{Dec}(\mathsf{sk}, C = (c,V))$ runs $s \leftarrow \mathsf{KEM}.\mathsf{Dec}(\mathsf{sk},c)$, gets $U'\|V' \leftarrow \mathsf{PRG}(s)$, and checks whether $V = V'$. In the positive case, one outputs $K' \leftarrow U'$, otherwise one outputs $\bot$.

Correctness. If the KEM KEM is correct, then the derived KEM' with authentication is also correct, has the decapsulation of c outputs the same s as during encapsulation, and then $\mathsf{PRG}(s)$ gives the same key and tag.

Security Properties. We will now show the previous security notions still hold, and we really provide authentication. We can claim that the above KEM' retains the initial security properties of the KEM scheme, but as the proofs essentially rely of the PRG properties, we defer the proofs to the Appendix D.

Theorem 6 (Session-Key Privacy). *If the KEM KEM is SK-IND-secure, and outputs (C,K)'s of $\mathsf{KEM}.\mathsf{Enc}$ have uniformly distributed K's in $\{0,1\}^\kappa$, then its derived KEM' with authentication using the IND-secure PRG $\mathsf{PRG} : \{0;1\}^\kappa \to \{0;1\}^{k+\ell}$ is SK-IND-secure:* $\mathsf{Adv}^{sk\text{-}ind}_{\mathsf{KEM}'}(\tau) \leq 2 \cdot \mathsf{Adv}^{sk\text{-}ind}_{\mathsf{KEM}}(\tau) + 2 \cdot \mathsf{Adv}^{ind}_{\mathsf{PRG}_{\kappa,k+\ell}}(\tau)$, *for any running time τ.*

Theorem 7 (Public-Key Privacy). *If the KEM KEM is both SK-IND and PK-IND-secure, outputs (C,K)'s of $\mathsf{KEM}.\mathsf{Enc}$ have uniformly distributed K's in $\{0,1\}^\kappa$, and $\mathsf{PRG} : \{0;1\}^\kappa \to \{0;1\}^{k+\ell}$ is an IND-secure PRG, then its derived KEM' using PRG is PK-IND-secure:* $\mathsf{Adv}^{pk\text{-}ind}_{\mathsf{KEM}'}(\tau) \leq \mathsf{Adv}^{pk\text{-}ind}_{\mathsf{KEM}}(\tau) + 4 \cdot \mathsf{Adv}^{sk\text{-}ind}_{\mathsf{KEM}}(\tau) + 4 \cdot \mathsf{Adv}^{ind}_{\mathsf{PRG}_{\kappa,k+\ell}}(\tau)$, *for any running time τ.*

We develop the authentication property, with the proof in the Appendix B:

Theorem 8 (Authentication). *If the KEM KEM is SK-IND, outputs (C,K)'s of $\mathsf{KEM}.\mathsf{Enc}$ have uniformly distributed K's in $\{0,1\}^\kappa$, and $\mathsf{PRG} : \{0;1\}^\kappa \to \{0;1\}^{k+\ell}$ is an IND-secure PRG, then the corresponding authenticated KEM KEM' using PRG provides authentication:* $\mathsf{Adv}^{auth}_{\mathsf{KEM}'}(\kappa) \leq 2^{-\ell} + \mathsf{Adv}^{sk\text{-}ind}_{\mathsf{KEM}}(\tau) + \mathsf{Adv}^{ind}_{\mathsf{PRG}_{\kappa,k+\ell}}(\tau')$, *for some small running times τ, τ'.*

4 Subset-Cover KEMAC

The above notion of access control is quite general and includes both key-policy ABE and ciphertext-policy ABE, where one can have policies $\mathcal{P}$ and attributes such that given a subset of attributes, this defines a list of Boolean B (according to the presence or not of the attribute), and $\mathcal{P}(B)$ is either true or false.

For efficiency considerations, we will focus on the subset-cover approach: during the Setup, one defines multiple sets S_i; when generating a user key usk_j, a list A_j of subsets if specified, which implicitly means user $U_j \in S_i$ for all $i \in A_j$; at encapsulation time, a target set T is given by B, such that $T = \cup_{i\in B}S_i$.

Intuitively, S_i's are subsets of the universe of users, and to specify the receivers, one encapsulates the key K for a covering of the target set T. A KEMAC, for a list Σ of sets S_i, can then be defined from any KEM in $\mathcal{K}$ that is a group with internal law denoted $\oplus$. The most basic version of such a KEMAC is provided in appendix E. In this section, we directly describe a subset cover KEMAC with anonymity and early aborts, as these will be the properties used in our final construction.

Anonymous Subset-Cover KEMAC with Early Aborts. To avoid sending B together with the ciphertext, but still being able to quickly find the correct matching indices in the ciphertext and the user's key, one can use a KEM' with authentication:

- $\mathsf{KEMAC}.\mathsf{Setup}(\Sigma)$, for each $S_i \in \Sigma$, runs $(\mathsf{pk}_i,\mathsf{sk}_i) \leftarrow \mathsf{KEM}'.\mathsf{KeyGen}(1^\kappa)$: $\mathsf{PK} \leftarrow (\mathsf{pk}_i)_i$ and $\mathsf{MSK} \leftarrow (\mathsf{sk}_i)_i$;

- KEMAC.KeyGen(MSK, A_j) defines the user's secret key $\mathsf{usk}_j \leftarrow (\mathsf{sk}_i)_{i \in A_j}$;
- KEMAC.Enc(PK, B) generates a random session key $K \xleftarrow{\$} \{0;1\}^k$, and, for all $i \in B$, runs $(C_i, K_i) \leftarrow$ KEM'.Enc(pk_i) and outputs $C \leftarrow (C_i, E_i = K \oplus K_i)_{i \in B}$ together with the encapsulated key K;
- KEMAC.Dec(usk, C), for all sk_i in usk and all (C_j, E_j) in C, runs $K'_{i,j} \leftarrow$ KEM'.Dec(sk_i, C_j). It stops for the first valid $K'_{i,j}$, outputs $K \leftarrow K'_{i,j} \oplus E_j$.

For this above scheme, we can claim both the SK-IND security and the AC-IND security, for selective key queries. But first, let us check the correctness, thats fails if a wrong key, among the $S_A S_B$ possibilities, makes accepts:

Theorem 9 (Correctness). *If the underlying KEM' is AUTH-secure, the above subset-cover KEMAC is correct:* $\mathsf{Adv}^{cor}_{KEMAC}(\kappa) \leq S_A S_B \times \mathsf{Adv}^{auth}_{KEM'}(\kappa)$, *where S_A and S_B are the sizes of the user' sets of attributes and the number of subsets in the ciphertext, respectively.*

About SK-IND and AC-IND security, the proofs follow the classical hybrid technique, they are thus deferred to the Appendix D.3.

Theorem 10 (Session-Key Privacy). *If the underlying KEM' is SK-IND-secure, the above subset-cover KEMAC is also SK-IND-secure, for selective key-queries:* $\mathsf{Adv}^{sk\text{-}ind}_{KEMAC}(\tau) \leq 2q_k \times \mathsf{Adv}^{sk\text{-}ind}_{KEM'}(\tau)$, *where q_k is the number of key-queries.*

Theorem 11 (Access-Control Privacy). *If the underlying KEM' is AC-IND-secure, the above subset-cover KEMAC is AC-IND-secure, for selective key-queries and constant-size sets B:* $\mathsf{Adv}^{ac\text{-}ind}_{KEMAC}(\tau) \leq 2S_B \times \mathsf{Adv}^{pk\text{-}ind}_{KEM}(\tau)$, *where S_B is the constant-size of the sets B.*

We stress that B must have a constant size to achieve access-control privacy.

5 Traceable KEM

In a subset-cover-based KEMAC, a same decapsulation key sk_i is given to multiple users, for a public key pk_i. In case of abuse, one cannot trace the defrauder. We offer an ElGamal-based KEM with traceability, in the same vein as [BF99].

Traceable ElGamal-based TKEM. Let $\mathbb{G}$ be a group of prime order q, with a generator g, in which the Computational Diffie-Hellman problem is hard. We describe below a TKEM with n multiple decapsulation keys for a specific public key, allowing to deal with collusions of at most t users:

- TKEM.KeyGen($1^\kappa, n, t, g, \mathbb{G}, q$): outputs a public key pk and n secret keys usk_j:
  - it samples random $s, s_k \xleftarrow{\$} \mathbb{Z}_q^*$, for $k = 1 \ldots, t+1$ and sets $h \leftarrow g^s$ as well as $h_k \leftarrow g^{s_k}$ for each k;
  - for users U_j, for $j = 1 \ldots, n$, one samples random $(v_{j,k})_k \xleftarrow{\$} \mathbb{Z}_q^{t+1}$, such that $\sum_k v_{j,k} s_k = s$, for $j = 1 \ldots, n$. Then, $\mathsf{pk} \leftarrow ((h_k)_k, h)$, while each $\mathsf{usk}_j \leftarrow (v_{j,k})_k$.
- TKEM.Enc($\mathsf{pk} = ((h_k)_k, h)$): it samples a random $r \xleftarrow{\$} \mathbb{Z}_q$, and sets $C = (C_k \leftarrow h_k^r)_k$, as well as $K \leftarrow h^r$.
- TKEM.Dec($\mathsf{usk}_j = (v_{j,k})_k, C = (C_k)_k$): it outputs $K \leftarrow \prod_k C_k^{v_{j,k}}$

One notes $\prod_k C_k^{v_{j,k}} = \prod_k h_k^{r v_{j,k}} = \prod_k (g^r)^{s_k v_{j,k}} = g^{r \sum_k s_k v_{j,k}} = g^{sr} = h^r = K$.

Security Properties. First, we will show that the above TKEM construction achieves both SK-IND and PK-IND security. But it also allows to confirm traitors, from a stateless pirate decoder $\mathcal{P}$ (in particular, this means that $\mathcal{P}$ never blocks itself after several invalid ciphertexts). The proofs of Theorems 12 and 13 can be found in appendices D.6 and D.7.

Theorem 12 (Session-Key Privacy). *The above TKEM achieves SK-IND security under the DDH assumption in $\mathbb{G}$:* $\mathsf{Adv}^{sk\text{-}ind}_{TKEM}(\tau) \leq \mathsf{Adv}^{ddh}_{\mathbb{G}}(\tau)$.

Theorem 13 (Public-Key Privacy). *The above TKEM achieves PK-IND security under the DDH assumption in $\mathbb{G}$:* $\mathsf{Adv}^{pk\text{-}ind}_{TKEM}(\tau) \leq \mathsf{Adv}^{ddh}_{\mathbb{G}}(\tau)$.

Theorem 14 (t-Confirmation). *A collusion of at most t keys can be confirmed from a useful stateless pirate decoder $\mathcal{P}$: starting from a correct guess for $\mathcal{T}$, the traitors' keys used for building the pirate decoder $\mathcal{P}$, by accessing the decoder, one can confirm a traitor in $\mathcal{T}$, with negligible error.*

Proof. To prove this theorem, we first give a description of the confirmer algorithm $\mathcal{C}$, then we provide the indistinguishability analysis, and eventually prove $\mathcal{C}$ will give a correct answer. This proof can be found in the Appendix C.

Corrolary 1 *In the particular case of $t = 1$, one can efficiently trace one traitor, from a useful stateless pirate decoder: by trying $\mathcal{G} = \{J\}$ sequentially for each $J = 1, \ldots, n$, and evaluating $p_\mathcal{G}$, one should get either a significant advantage (for the traitor) or 0 (for honest keys).*

6   Our **KEMAC** Scheme

We have already presented a traceable KEM that is secure against classical adversaries. If we combine it with another scheme expected secure against quantum adversaries, we can thereafter combine them into an hybrid-KEM, that inherits security properties from both schemes, with still traceability against classical adversaries. But we will actually exploit the properties of a Public-Key Encryption (PKE) scheme in order to improve efficiency of the combination. Given a PKE, that is both indistinguishable and anonymous, we can trivially get a KEM that is both SK-IND and PK-IND secure:

- KEM.KeyGen(1^κ) gets $(\mathsf{pk}, \mathsf{sk}) \leftarrow$ PKE.KeyGen(1^κ), and outputs $(\mathsf{pk}, \mathsf{sk})$;
- KEM.Enc(pk) gets $K \xleftarrow{\$} \mathcal{K}$, $C \leftarrow$ PKE.Enc(pk, K), and outputs (K, C);
- KEM.Dec(sk, C) outputs PKE.Dec(sk, C).

CRYSTALS-Kyber PKE We recall the algorithms of the CRYSTALS-Kyber [ABD+21] public-key encryption whose both indistinguishability and anonymity rely on the hardness of Module-LWE [LS15]. We identify R_q with $\mathbb{Z}_q^n$ that contains the plaintext space $\mathcal{K} = \{0; 1\}^n$, and use two noise parameters $\eta_1 \geq \eta_2$, for the Gaussian distributions $\mathcal{B}_{\eta_1}$ and $\mathcal{B}_{\eta_2}$:

- Kyber.KeyGen(1^κ): sample random $\mathbf{A} \xleftarrow{\$} \mathsf{R}_q^{k \times k}$ and $(\mathbf{s}, \mathbf{e}) \xleftarrow{\$} \mathcal{B}_{\eta_1}^k \times \mathcal{B}_{\eta_1}^k$, then set $\mathsf{pk} \leftarrow (\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ and $\mathsf{sk} \leftarrow \mathbf{s}$.
- Kyber.Enc(pk, K): $\mathbf{r} \xleftarrow{\$} \mathcal{B}_{\eta_1}^k$, and $(\mathbf{e}_1, e_2) \xleftarrow{\$} \mathcal{B}_{\eta_2}^k \times \mathcal{B}_{\eta_2}$, then set $\mathbf{u} = \mathbf{A}^T \mathbf{r} + \mathbf{e}_1$ and $v = \mathbf{b}^T \mathbf{r} + e_2 + \lceil \frac{q}{2} \rfloor \cdot K$, and return $C = (\mathbf{u}, v)$.
- Kyber.Dec(sk, C): compute $w \leftarrow v - \mathbf{s}^T \mathbf{u}$ and output $K = \lceil \frac{2}{q} \cdot w \rfloor$.

In the Appendix G, we recall the correctness, and the security proofs for indistinguishability and anonymity of this scheme, assuming $\eta_1 \geq \eta_2$ (while in practice, we often have $\eta_1 = \eta_2$). Theorem 15 follows from [ABD+21], and Theorem 16 is also in the scope of [MX22]:

Theorem 15 (Indistinguishability of Kyber.). Kyber *is IND-secure under the decisional Module-LWE assumption:*
$$\mathsf{Adv}_{\mathsf{Kyber}}^{\mathsf{ind}}(\tau) \leq \mathsf{Adv}_{\mathsf{R}_q, k, k, \eta_1}^{\mathsf{dmlwe}}(\tau) + \mathsf{Adv}_{\mathsf{R}_q, k+1, k, \eta_2}^{\mathsf{dmlwe}}(\tau) \leq 2 \times \mathsf{Adv}_{\mathsf{R}_q, k+1, k, \eta_2}^{\mathsf{dmlwe}}(\tau).$$

Theorem 16 (Anonymity of Kyber.). Kyber *is PK-IND-secure under the decisional Module-LWE assumption:*
$$\mathsf{Adv}_{\mathsf{Kyber}}^{pk\text{-}ind}(\tau) \leq 2 \times \mathsf{Adv}_{\mathsf{R}_q, k, k, \eta_1}^{\mathsf{dmlwe}}(\tau) + \mathsf{Adv}_{\mathsf{R}_q, k+1, k, \eta_2}^{\mathsf{dmlwe}}(\tau) \leq 3 \times \mathsf{Adv}_{\mathsf{R}_q, k+1, k, \eta_2}^{\mathsf{dmlwe}}(\tau).$$

Hybrid KEM, from KEM and PKE. Using the ElGamal KEM that is both SK-IND and PK-IND-secure under the DDH assumption, together with the Kyber PKE that is both SK-IND and PK-IND-secure under the DMLWE assumption, the hybrid KEM is:

- SK-IND-secure, as soon as either the DDH or the DMLWE assumptions hold;
- PK-IND-secure, under both the DDH and the DMLWE assumption.

according to Section 2. But with a PKE scheme, we can optimize a bit with:

- Hyb.KeyGen(1^κ): generate both pairs of keys $(\mathsf{pk}_1, \mathsf{sk}_1) \leftarrow$ KEM.KeyGen(1^κ) and $(\mathsf{pk}_2, \mathsf{sk}_2) \leftarrow$ PKE.KeyGen(1^κ), then output $\mathsf{pk} \leftarrow (\mathsf{pk}_1, \mathsf{pk}_2)$ and $\mathsf{sk} \leftarrow (\mathsf{sk}_1, \mathsf{sk}_2)$;
- Hyb.Enc(pk): parse pk as $(\mathsf{pk}_1, \mathsf{pk}_2)$, choose a random $K \xleftarrow{\$} \mathcal{K}$, call $(C_1, K_1) \leftarrow$ KEM.Enc(pk_1) and $C_2 \leftarrow$ PKE.Enc($\mathsf{pk}_2, K \oplus K_1$). Output $(C = (C_1, C_2), K)$;
- Hyb.Dec(sk, C): parse sk as $(\mathsf{sk}_1, \mathsf{sk}_2)$ and C as (C_1, C_2), then call both $K_1 \leftarrow$ KEM.Dec(sk_1, C_1), $K_2 \leftarrow$ PKE.Dec(sk_2, C_2), and output $K = K_1 \oplus K_2$.

Hybrid Traceable KEMAC. We can apply the above generic combination to build an anonymous subset-cover KEMAC with early abort, with the traceable ElGamal KEM and Kyber PKE to get a Key Encapsulation Mechanism with Access Control and Black-Box traceability (without collusions, so with $t = 1$ using notations from Section 5), where all the privacy notions (message-privacy and target-set privacy) hold as soon as at least the DDH or the DMLWE assumptions hold, while traceability works under the DDH assumption.

To have authentication properties, the ElGamal TKEM is slightly modified to fit theorems 6, 7 and 8's requirements, in which the element K output by the encapsulation algorithm should be uniform in $\{0;1\}^\kappa$. This modification can be done either in the Random Oracle Model (ROM) with a hash function modelled as a random oracle, and outputting a hash of the original key into $\{0;1\}^\kappa$, or, without the ROM, using a twist augmented technique from [CFGP06]. The KEMs derived with these two techniques are presented in appendix F, along with their security proofs. We describe here the one in the ROM. Proofs for SK-IND and PK-IND-securities follow immediately from the proofs that TKEM is SK-IND and PK-IND-secure.

Detailed Description. The straightforward construction of the hybrid traceable KEMAC with early abort is the simple instantiation of the KEMAC scheme from Section 4 from a KEM with authentication (from Section 3), itself based on our hybrid KEM from the previous subsection. A naïve instantiation would draw independent keys in the hybrid schemes and send their $\oplus$'s with the encapsulated key. But as K is chosen beforehand, the same K can be chosen for all the subsets. This optimized version is described with the following algorithms, where $\mathcal{H}$ is a hash function modeled as a random oracle with output length κ[5], PRG : $\{0;1\}^\kappa \to \{0;1\}^{k+\ell}$ a PRG, where k is the length of the encapsulated key, ℓ the length of the verification tag, and Σ the set of subsets $(S_i)_i$ (or attributes). We instantiate it with the Kyber PKE, but it would work with any PKE that is both indistinguishable and anonymous. We call this KEMAC Covercrypt:

- Covercrypt.Setup($\Sigma, 1^\kappa$):
  1. For a group $\mathbb{G}$ of prime order p, generated by g, one samples $s, s_1, s_2 \stackrel{\$}{\leftarrow} \mathbb{Z}_p$, then sets $h = g^s$, and $g_1 = g^{s_1}, g_2 = g^{s_2}$ (for tracing purposes).
  2. Then, for tracing, we set $\mathsf{tsk} = (s, s_1, s_2, \mathcal{ID})$, where $\mathcal{ID}$ is the set of the users' identifiers uid, initialized as an empty set here, and $\mathsf{tpk} = (g, h, g_1, g_2)$.
  3. For each $S_i \in \Sigma$, one samples a random scalar $x_i \stackrel{\$}{\leftarrow} \mathbb{Z}_p$, a $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{Kyber.KeyGen}(1^\kappa)$, then sets $\mathsf{pk}'_i \leftarrow (h_i = g^{x_i}, \mathsf{pk}_i)$, and $\mathsf{sk}'_i \leftarrow (x_i, \mathsf{sk}_i)$.
  4. Finally, the global public key is set to $\mathsf{PK} \leftarrow (\mathsf{tpk}, \{\mathsf{pk}'_i\}_i)$, and the master secret key to $\mathsf{MSK} \leftarrow (\mathsf{tsk}, \{\mathsf{sk}'_i\}_i, \mathcal{UP})$, where $\mathcal{UP}$ is the set of user's secret keys, showing their permissions, but initialized as an empty set. One returns $(\mathsf{MSK}, \mathsf{PK})$.
- Covercrypt.KeyGen(MSK, U, A):
  1. For a user U, with attributes A (a list of subsets, or equivalently their indices), one samples $(\alpha, \beta) \in \mathbb{Z}_p^2$ such that $\alpha s_1 + \beta s_2 = s$, and sets the corresponding user secret identifier $\mathsf{uid} \leftarrow (\alpha, \beta)$.
  2. The tracing secret key tsk is updated as tsk' by adding (U, uid) in $\mathcal{ID}$.
  3. Finally, the user's secret key is defined as $\mathsf{usk} \leftarrow (\mathsf{uid}, \{\mathsf{sk}'_j\}_{j\in A})$, and one outputs it along with MSK', the master secret key MSK updated with usk added in $\mathcal{UP}$, and tsk' instead of tsk.
- Covercrypt.Enc(PK, B):
  1. For a target set that covers all the users with an attribute in B (or equivalently the indices of attributes, such that $A \cap B \neq \emptyset$), one generates a random seed for the key to be encapsulated, $S \stackrel{\$}{\leftarrow} \{0;1\}^\kappa$, then draws $r \stackrel{\$}{\leftarrow} \mathbb{Z}_p^*$, sets $c = (C_1 = g_1^r, C_2 = g_2^r)$, and, for each $i \in B$, with $\mathsf{pk}'_i = (h_i = g^{x_i}, \mathsf{pk}_i)$, sets $K_i = \mathcal{H}(h_i^r)$, and then sets $E_i \leftarrow \mathsf{Kyber.Enc}(\mathsf{pk}_i, S \oplus K_i)$[6].
  2. One then computes $K||V \leftarrow \mathsf{PRG}(S)$, in order to grant the early aborts paradigm, and sets the encapsulation as: $C \leftarrow (c, \{E_i\}_{i\in B}, V)$, the encapsulated key as K, and outputs: (K, C).
- Covercrypt.Dec($\mathsf{usk} = (\mathsf{uid} = (\alpha, \beta), \{\mathsf{sk}'_j\}_{j\in A}), C = (c, \{E_i\}_{i\in B}, V)$): For each $\mathsf{sk}'_j = (x_j, \mathsf{sk}_j)$ in usk and (c, E_i, V) in C (for one index i), one decapsulates the underlying hybrid KEM to get the potential seed S used for the key:
  - first, $K'_{i,j} \leftarrow \mathsf{Kyber.Dec}(\mathsf{sk}_j, E_i)$;
  - for ElGamal, from $c = (C_1, C_2)$, one computes $K_j \leftarrow \mathcal{H}((C_1^\alpha C_2^\beta)^{x_j})$;

---

[5] Our base pre-quantum KEM is thus the ROM.UTKEM described in appendix F.

[6] Note that this is the optimized version of a generic one where one would have drawn $|B|$ extra session keys K'_i, E_i would actually have been a Kyber encryption of these K'_i's instead of the $S \oplus K_i$, and one would have had to send $|B|$ extra $F_i \leftarrow K_i \oplus K' \oplus S$.

- $S_{i,j}$ is then computed as $S_{i,j} \leftarrow K'_{i,j} \oplus K_j$.

In the early-abort check, one computes $U'_{i,j} || V'_{i,j} \leftarrow \mathsf{PRG}(S_{i,j})$, and checks whether $V'_{i,j} = V$. In the positive case, one returns $K \leftarrow U'_{i,j}$, for this first valid (i,j), as the session key. Else, if $V'_{i,j} \neq V$, the ciphertext is rejected and the loop on the i,j indices goes on[7].

Security Analysis. Our Covercrypt scheme inherits its security properties from the underlying hybrid KEM scheme using both the Kyber PKE and the traceable ElGamal KEM, and as such, is SK-IND-secure as soon as either the DDH or the DMLWE assumptions hold, and PK-IND-secure under both the DDH and the DMLWE assumptions. Correctness also follows from the authentication property of the hybrid KEM, and thus under either the DDH or the DMLWE assumptions.

Traceability. The traceability is inherited from the underlying traceable ElGamal KEM scheme, with $t = 1$ in Section 5's notations; it relies on the DDH. To check whether a user U with $\mathsf{uid} = (\alpha, \beta)$ using the key sk – which is shared among her and other users – is corrupted, one encapsulates a key that only this user can decapsulate with sk, because the ElGamal encapsulations are group elements with exponent a random linear combination of a vector which is orthogonal to (α, β), following the confirmer construction from Section 5. We stress that our construction with $t = 1$ does not allow collusions. But it can be extended to confirm larger t-big collusions of traitors.

7 Implementation

Parameters of Covercrypt. The parameters of Kyber are recalled in the Appendix G, Table 3, with the sizes (in Bytes) of the keys and ciphertexts, using the compression/decompression, as this does not impact our security results, and the Kyber PKE can be used as a black-box from any library.

We have done an implementation in Rust of Covercrypt (a pre- and post-quantum hybridized Anonymous Subset-Cover KEMAC with Early-Aborts), with optimization for a security of 128 bits[8]. We use Kyber-768 (and its `pqd_kyber` library[9]) and ElGamal on the Curve25519, as group that is of prime order $p = 2^{255} - 19$. The hash algorithm used to generate the Early-Abort tags (256 bits) and the keys (256 bits) generated by the KEM is SHAKE-256. Then we present the sizes of the keys and ciphertexts, according to the sizes of A and B, in Table 1. We compare these with the sizes obtained for a KEM based on a pre-quantum [GPSW06] ABE scheme[10], way more efficient than post-quantum ones such as [DDP+18] [11].

Size of A	1	2	3	4	5
Covercrypt Secret Key usk	1250	2435	3620	4805	5990
Coverc. Pre-Quant. S. K. (uid, $\{x_i\}_i$)	98	131	164	197	230
User Secret Key with GPSW	340	504	668	832	996
Size of B	1	2	3	4	5
Covercrypt Encapsulation C	1171	2260	3349	4438	5527
Covercrypt Pre-Quant. Encaps. (c, V)	115	148	181	214	247
GPSW KEM Encapsulation	400	452	504	556	608

Table 1. Sizes of keys and encapsulations (in Bytes) according the sizes of A and B.

---

[7] Again, this corresponds to our optimized version, taking advantage of the encrypting properties of Kyber. For a generic hybrid KEMAC, one would have output $U'_{i,j} \oplus F_i$ when $V'_{i,j} = V$ (cf. previous footnote for the definition of F_i).

[8] https://github.com/Cosmian/cover_crypt

[9] https://docs.rs/pqc_kyber/latest/pqc_kyber/

[10] Whose implementation can be found at: https://github.com/Cosmian/abe_gpsw.

[11] In this comparison, to translate the attribute setting into a subset-cover one, we consider a context in which users hold $|A| + 1$ attributes, corresponding to $|A|$ subsets in the subset-cover setting, the subsets being the intersection of one of these attributes with each of the other ones, and that encapsulations are made in the same way with respect to $|B| + 1$ attributes corresponding to $|B|$ subsets, and for the decapsulation timings, we suppose there is always exactly one subset in the intersection of the ones the user has access to and the ones in the encapsulation.

Benchmarks The benchmarks in table 2 are performed on an Intel Core Processor (Haswell, no TSX) CPU @3MHz. The table shows the time required to generate Covercrypt encapsulations and decapsulations for a 32-Byte symmetric key, with the same definitions for the sizes $|A|$ and $|B|$ as in table 1. These performances are, as before, compared with the [GPSW06]-based KEM's.

Size of B	1	2	3	4	5
Covercrypt	191	272	329	401	487
GPSW KEM	4793	5431	6170	6607	7245

Encapsulation time (in μs)

$|A|\downarrow \quad \backslash |B| \rightarrow$	1	2	3	4	5
1	214	247	288	345	454
2	311	386	466	543	562
3	334	400	505	608	702
4	471	613	781	908	1072
5	467	646	831	1058	1212

Covercrypt decapsulation time (in μs)

Table 2. Comparisons of Covercrypt and GPSW-based encapsulation/decapsulation times. For decapsulation, the GPSW-based KEM has a constant runtime of approximately 3880 μs.

References

ABD+21. Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-kyber algorithm specifications and supporting documentation. `https://pq-crystals.org/kyber/resources.shtml`, 2021.

App12. Benny Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 805–816. ACM Press, May 2012.

BBDP01. Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 566–582. Springer, Heidelberg, December 2001.

BBF+19. Nina Bindel, Jacqueline Brendel, Marc Fischlin, Brian Goncalves, and Douglas Stebila. Hybrid key encapsulation mechanisms and authenticated key exchange. In Jintai Ding and Rainer Steinwandt, editors, *Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019*, pages 206–226. Springer, Heidelberg, 2019.

BF99. Dan Boneh and Matthew K. Franklin. An efficient public key traitor tracing scheme. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 338–353. Springer, Heidelberg, August 1999.

CFGP06. Olivier Chevassut, Pierre-Alain Fouque, Pierrick Gaudry, and David Pointcheval. The Twist-AUgmented technique for key exchange. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006*, volume 3958 of *LNCS*, pages 410–426. Springer, Heidelberg, April 2006.

DDP+17. Wei Dai, Yarkın Doröz, Yuriy Polyakov, Kurt Rohloff, Hadi Sajjadpour, Erkay Savaş, and Berk Sunar. Implementation and evaluation of a lattice-based key-policy ABE scheme. Cryptology ePrint Archive, Report 2017/601, 2017. `https://eprint.iacr.org/2017/601`.

DDP+18. Wei Dai, Yarkin Doröz, Yuriy Polyakov, Kurt Rohloff, Hadi Sajjadpour, Erkay Savas, and Berk Sunar. Implementation and evaluation of a lattice-based key-policy ABE scheme. *IEEE Trans. Inf. Forensics Secur.*, 13(5):1169–1184, 2018.

FP12. Nelly Fazio and Irippuge Milinda Perera. Outsider-anonymous broadcast encryption with sublinear ciphertexts. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 225–242. Springer, Heidelberg, May 2012.

GHP18. Federico Giacon, Felix Heuer, and Bertram Poettering. KEM combiners. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part I*, volume 10769 of *LNCS*, pages 190–218. Springer, Heidelberg, March 2018.

GPSW06. Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 89–98. ACM Press, October / November 2006. Available as Cryptology ePrint Archive Report 2006/309.

HILL99.   Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

LG18.   Jiangtao Li and Junqing Gong. Improved anonymous broadcast encryptions - tight security and shorter ciphertext. In Bart Preneel and Frederik Vercauteren, editors, *ACNS 18*, volume 10892 of *LNCS*, pages 497–515. Springer, Heidelberg, July 2018.

LPQ12.   Benoît Libert, Kenneth G. Paterson, and Elizabeth A. Quaglia. Anonymous broadcast encryption: Adaptive security and efficient constructions in the standard model. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 206–224. Springer, Heidelberg, May 2012.

LS15.   Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.*, 75(3):565–599, 2015.

MX22.   Varun Maram and Keita Xagawa. Post-quantum anonymity of kyber. Cryptology ePrint Archive, Report 2022/1696, 2022. `https://eprint.iacr.org/2022/1696`.

Sho01.   Victor Shoup. A proposal for an iso standard for public key encryption. `https://shoup.net/papers/iso-2_1.pdf`, December 2001.

Wee21.   Hoeteck Wee. ABE for DFA from LWE against bounded collusions, revisited. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part II*, volume 13043 of *LNCS*, pages 288–309. Springer, Heidelberg, November 2021.

Appendix

A Public-Key Encryption

A Public-Key Encryption (PKE) scheme is defined by 3 algorithms:

- $\mathsf{PKE.KeyGen}(1^\kappa)$: the *key generation algorithm* outputs a pair of public and secret keys $(\mathsf{pk}, \mathsf{sk})$;
- $\mathsf{PKE.Enc}(\mathsf{pk}, m)$: the *encryption algorithm* encrypts the input message m under the public key pk and outputs the ciphertext C;
- $\mathsf{PKE.Dec}(\mathsf{sk}, C)$: the *decryption algorithm* outputs the message m encrypted in C.

We will use the classical notion of indistinguishability and of anonymity of such a PKE scheme, similarly to the same notions for KEMs:

Indistinguishability. A PKE is *indistinguishable* (denoted IND) if for an honestly generated pk, when the adversary chooses two messages m_0 and m_1, it cannot distinguish an encryption of m_0 from an encryption of m_1, both under pk. More formally, a PKE is IND-secure if for any adversary $\mathcal{A}$, $\mathsf{Adv}_{\mathsf{PKE}}^{\mathsf{ind}}(\mathcal{A}) = \mathsf{negl}(\kappa)$, for $b \xleftarrow{\$} \{0; 1\}$ and

$$\mathcal{D}_b = \left\{ \begin{array}{l} (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{PKE.KeyGen}(1^\kappa), \\ (m_0, m_1, s) \leftarrow \mathcal{A}(\mathsf{pk}), C \leftarrow \mathsf{PKE.Enc}(\mathsf{pk}, m_b) \end{array} : (\mathsf{pk}, C, s) \right\}$$
$$\mathsf{Adv}_{\mathsf{PKE}}^{\mathsf{ind}}(\mathcal{A}) = 2 \times \Pr_{\mathcal{D}_b}[\mathcal{A}(\mathsf{pk}, C, s) = b] - 1.$$

where s is an internal state the adversary keeps between the two steps.

Anonymity. A PKE is *anonymous*, also known as *public-key privacy* (denoted $\mathsf{PK\text{-}IND}$) if for two honestly generated pk_0 and pk_1, when the adversary chooses a message m, it cannot distinguish an encryption of m under pk_0 from the encryption of m under pk_1. More formally, a PKE is $\mathsf{PK\text{-}IND}$-secure if for any adversary $\mathcal{A}$, $\mathsf{Adv}_{\mathsf{PKE}}^{\mathsf{pk\text{-}ind}}(\mathcal{A}) = \mathsf{negl}(\kappa)$, for $b \xleftarrow{\$} \{0; 1\}$ and

$$\mathcal{D}_b = \left\{ \begin{array}{l} \text{For } i = 0, 1: \\ \quad (\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{PKE.KeyGen}(1^\kappa), \\ (m, s) \leftarrow \mathcal{A}(\mathsf{pk}_0, \mathsf{pk}_1), C \leftarrow \mathsf{PKE.Enc}(\mathsf{pk}_b, m) \end{array} : (\mathsf{pk}_0, \mathsf{pk}_1, C, s) \right\}$$
$$\mathsf{Adv}_{\mathsf{PKE}}^{\mathsf{pk\text{-}ind}}(\mathcal{A}) = 2 \times \Pr_{\mathcal{D}_b}[\mathcal{A}(\mathsf{pk}_0, \mathsf{pk}_1, C, s) = b] - 1.$$

B Proof of Theorem 8

We present a sequence of games, from the AUTH security game against KEM'.

Game $\mathbf{G_0}$: In the initial game, one runs $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KEM}'.\mathsf{KeyGen}(1^\kappa)$, $(c, s) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk}_0)$ and $K_0 \| V \leftarrow \mathsf{PRG}(s)$. One then runs $s' \leftarrow \mathsf{KEM.Dec}(\mathsf{sk}_1, c)$, followed by $U' \| V' \leftarrow \mathsf{PRG}(s')$. We denote P_0 the probability $V' = V$. This is $\mathsf{Adv}_{\mathsf{KEM}'}^{\mathsf{auth}}(1^\kappa)$.

Game $\mathbf{G_1}$: In this game, we just replace $s \xleftarrow{\$} \{0; 1\}^\kappa$, that is drawn uniformly at random from the session-key space of KEM, $\{0; 1\}^\kappa$. The difference between this game and the previous one is the $\mathsf{SK\text{-}IND}$-game on the underlying KEM, against a trivial adversary $\mathcal{A}_0$. Hence, $P_0 - P_1 \leq \mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{sk\text{-}ind}}(\tau)$, τ the running time of the trivial adversary $\mathcal{A}_0$ that runs two key generations, one encapsulation, two PRG evaluations, and one decapsulation.

Game $\mathbf{G_2}$: In this game, one takes $K_0 \| V \xleftarrow{\$} \{0; 1\}^{k+\ell}$. This is indistinguishable from the previous game except with probability $\mathsf{Adv}_{\mathsf{PRG}_{\kappa, k+\ell}}^{\mathsf{ind}}(\tau')$. Hence, $P_1 - P_2 \leq \mathsf{Adv}_{\mathsf{PRG}_{\kappa, k+\ell}}^{\mathsf{ind}}(\tau')$, where τ' is the running time of another trivial adversary $\mathcal{A}_1$ that runs two key generations, one encapsulation, one PRG evaluations, and one decapsulation.

In this game, as V is drawn uniformly at random from $\{0; 1\}^\ell$, the probability that it is equal to $V' \in \{0; 1\}^\ell$ is equal to $2^{-\ell}$: $P_2 = 2^{-\ell}$.

Finally, from the above, one deducts that:

$$\mathsf{Adv}_{\mathsf{KEM}'}^{\mathsf{auth}}(\kappa) \leq 2^{-\ell} + \mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{sk\text{-}ind}}(\tau) + \mathsf{Adv}_{\mathsf{PRG}_{\kappa, k+\ell}}^{\mathsf{ind}}(\tau')$$

C   Proof of Theorem 14

To prove this theorem, we first give a description of the confirmer algorithm $\mathcal{C}$, then we provide the indistinguishability analysis, and eventually prove $\mathcal{C}$ will give a correct answer.

Description of the Confirmer $\mathcal{C}$: The confirmer algorithm $\mathcal{C}$ can proceed as follows, for a candidate subset $\mathcal{G}$: $\{\mathsf{usk}_j = (v_{j,k})_k\}_{j \in \mathcal{G}}$, for $\mathcal{G}$ of size at most t: it chooses $(u_k)_k$ orthogonal to the subvector-space spanned by $\{(v_{j,k})_k\}_{j \in \mathcal{G}}$, which means that: $\sum_k u_k v_{j,k} = 0, \forall j \in \mathcal{G}$. This is possible as $(v_{j,k})_{k \in [1,t+1], j \in \mathcal{G}}$ is of rank at most t in $\mathbb{Z}_q^{t+1}$. Then the kernel is of dimension at least 1. One generates a fake ciphertext $C = (C_k)_k$, with $C_k \leftarrow h_k^r \cdot g^{u_k s'}$, for random $r, s' \xleftarrow{\$} \mathbb{Z}_q$, and then $K \leftarrow h^r$:

- Any key usk_j in $\mathcal{G}$ will lead to:

$$\prod_k C_k^{v_{j,k}} = \prod_k g^{(r s_k + s' u_k) \cdot v_{j,k}} = g^{r \sum_k s_k v_{j,k} + s' \sum_k u_k v_{j,k}} = g^{rs + s' \times 0} = K;$$

- and any key usk_j outside $\mathcal{G}$ will lead to: $\prod_k C_k^{v_{j,k}} = K \times (g^{\sum_k u_k v_{j,k}})^{s'} \neq K.$

we will show this allows to confirm at least one traitor from a candidat subset of traitors.

Indistinguishability Analysis. The above remark about the output key from a pirate decoder $\mathcal{P}$ assumes an honest behavior, whereas it can stop answering if it detects the fake ciphertext. We first need to show that, with the public key $\mathsf{pk} = ((h_k)_k, h)$ and only $\{\mathsf{usk}_j = (v_{j,k})_k\}_{j \in \mathcal{G}}$, one cannot distinguish the fake ciphertext from a real ciphertext, generated as above: from a Diffie-Hellman tuple $(A = g^a, B = g^r, C)$, one can derive, from random scalars $s, s'_k, u_k \xleftarrow{\$} \mathbb{Z}_q$, such that $\sum_k v_{j,k} s'_k = s$ and $\sum_k v_{j,k} u_k = 0$, for $j = 1 \ldots, n$:

$$h_k \leftarrow A^{u_k} \cdot g^{s'_k} = g^{a u_k + s'_k} \qquad\qquad h \leftarrow g^s \qquad\qquad \mathsf{usk}_j = (v_{j,k})_k \text{ for } j \in \mathcal{G}$$

where we implicitly define $s_k \leftarrow a u_k + s'_k$, that satisfy

$$\sum_k v_{j,k} s_k = \sum_k v_{j,k}(s'_k + a u_k) = \sum_k v_{j,k} s'_k + a \sum_k v_{j,k} u_k = s + 0 = s.$$

Then, one defines $C_k \leftarrow C^{u_k} \cdot B^{s'_k}$ and $K \leftarrow B^s$.

Let us note $C = g^{r-c}$, where c is either 0 (a Diffie-Hellman tuple) or random:

$$C_k = A^{(r+c)u_k} \cdot g^{r s'_k} = (A^{u_k} \cdot g^{s'_k})^r \cdot A^{c u_k} = h_k^r \cdot (A^c)^{u_k}.$$

One can remark that: when $c = 0$ (Diffie-Hellman tuple), $C = (C_k)_k$ is a normal ciphertext; when $c = s'$ (random tuple), this is a fake ciphertext. Under the DDH assumption, they are thus indistinguishable for an adversary knowing the keys $(\mathsf{usk}_i)_{i \in \mathcal{G}}$.

Confirmation of a Traitor. The above analysis shows that a pirate decoder $\mathcal{P}$ built from $(\mathsf{usk}_i)_{i \in \mathcal{G}}$ cannot distinguish the fake ciphertext from a real ciphertext. A useful pirate decoder should necessarily distinguish real key from random key. Then, several situations may appear, according to the actual set $\mathcal{T}$ of traitors' keys used to build the pirate decoder $\mathcal{P}$ by the adversary $\mathcal{A}$:

- If $\mathcal{T} \subseteq \mathcal{G}$, a useful decoder $\mathcal{P}$ can distinguish keys;
- If $\mathcal{T} \cap \mathcal{G} = \emptyset$, $\mathcal{P}$ cannot distinguish keys, as it can get several candidates, independent from the real or random keys.

Let us now assume we started from $\mathcal{G} \supseteq \mathcal{T}$, then the advantage of $\mathcal{P}$ in distinguishing real and random keys, denoted $p_{\mathcal{G}}$, is non-negligible, from the usefulness of the decoder. The following steps would also work if one starts with $\mathcal{G} \cap \mathcal{T} \neq \emptyset$, so that the advantage $p_{\mathcal{G}}$ is significant.

One then removes a user J from $\mathcal{G}$ to generate $\mathcal{G}'$ and new ciphertexts to evaluate $p_{\mathcal{G}'}$: if $J \notin \mathcal{T}$, usk_J is not known to the adversary, and so there is no way to check whether $\sum_k v_{J,k} s'_k = s$ and $\sum_k v_{J,k} u_k = 0$, even for a powerful adversary. So necessarily, $p_{\mathcal{G}'} = p_{\mathcal{G}}$.

On the other hand, we know that $p_{\emptyset} = 0$. So, one can sequentially remove users until a significant gap appears: this is necessarily for a user in $\mathcal{T}$. $\qquad\qquad\square$

D Other Deferred Proofs

D.1 Proof of theorem 6

We present a sequence of games, starting from the initial SK-IND security game against KEM':

Game $\mathbf{G}_0$: In the initial game, a challenger runs $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KEM.KeyGen}(1^\kappa)$, $(c, s) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk})$, and evaluates $U\|V \leftarrow \mathsf{PRG}(s)$ to output $(\mathsf{pk}, C = (c, V), K_b)$, where $b \xleftarrow{\$} \{0; 1\}$ is a random bit, $K_0 = U$, and $K_1 \xleftarrow{\$} \{0; 1\}^k$. From this output, the adversary returns her guess b' for b. We denote P_0 the probability of event $b' = b$ happening, which is $(1 + \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}'}(\mathcal{A}))/2$.

Game $\mathbf{G}_1$: In this game, $s \xleftarrow{\$} \{0; 1\}^\kappa$ is this time drawn uniformly at random from the session-key space of KEM, $\{0; 1\}^\kappa$, and one sets: $U\|V \leftarrow \mathsf{PRG}(s)$. The difference is the SK-IND-game on the underlying KEM. Hence, $P_0 - P_1 \leq \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\tau)$, where τ is the maximum running-time of adversary $\mathcal{A}$.

Game $\mathbf{G}_2$: In this game, one takes $U\|V \xleftarrow{\$} \{0; 1\}^{k+\ell}$, which is indistinguishable unless from the previous game except with probability $\mathsf{Adv}^{\mathsf{ind}}_{\mathsf{PRG}_{\kappa, k+\ell}}(\tau)$, where τ is the maximum running time of adversary $\mathcal{A}$. Hence, $P_1 - P_2 \leq \mathsf{Adv}^{\mathsf{ind}}_{\mathsf{PRG}_{\kappa, k+\ell}}(\tau)$. In this final game, this is clear that $P_2 = 1/2$, as U and K are both randomly drawn in $\{0; 1\}^k$.

Hence, $\mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}'}(\mathcal{A}) \leq 2 \cdot \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\tau) + 2 \cdot \mathsf{Adv}^{\mathsf{ind}}_{\mathsf{PRG}_{\kappa, k+\ell}}(\tau)$. $\qquad\qquad\qquad\qquad\qquad\square$

D.2 Proof of theorem 7

We present a sequence of games, starting from the initial PK-IND security game against KEM':

Game $\mathbf{G}_0$: In this game, for every i in $\{0; 1\}$, a challenger runs $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KEM.KeyGen}(1^\kappa)$, $(c_i, s_i) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk}_i)$, and gets $U_i\|V_i \leftarrow \mathsf{PRG}(s_i)$, then outputs $(\mathsf{pk}_0, \mathsf{pk}_1, C_b = (c_b, V_b))$, for a random $b \xleftarrow{\$} \{0; 1\}$. From the output, an adversary $\mathcal{A}$ outputs her guess b' for b. We denote P_0 the probability of the event $b' = b$ happening, which is equal to $(1 + \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}'}(\mathcal{A}))/2$.

Game $\mathbf{G}_1$: This game is as the previous one, except that now, one uses $s_i \xleftarrow{\$} \{0; 1\}^\kappa$ to evaluate $U_i\|V_i \leftarrow \mathsf{PRG}(s_i)$, for every i in $\{0; 1\}$. The correct guess probability difference is the SK-IND-game on the underlying KEM, for both i's in $\{0; 1\}$. Hence, $P_0 - P_1 \leq 2 \times \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\tau)$, where τ is the maximum running-time of adversary $\mathcal{A}$.

Game $\mathbf{G}_2$: This game is as the previous one, except that now, for each $i \in \{0; 1\}$, the challenger draws $U_i\|V_i \xleftarrow{\$} \{0; 1\}^{k+\ell}$. This is indistinguishable from the previous game under the security of $\mathsf{PRG} : \{0; 1\}^\kappa \rightarrow \{0; 1\}^{k+\ell}$. Hence, $P_1 - P_2 \leq \mathsf{Adv}^{\mathsf{ind}}_{\mathsf{PRG}_{\kappa, k+\ell}}(\tau)$, where τ is the maximal running time of the adversary $\mathcal{A}$.

Game $\mathbf{G}_3$: Now, one outputs $(\mathsf{pk}_0, \mathsf{pk}_1, C_b = (c_b, V))$, for a random $b \xleftarrow{\$} \{0; 1\}$ and a random $V \xleftarrow{\$} \{0; 1\}^\ell$. This simulation is perfectly indistinguishable from the previous game: $P_2 = P_3$. And this final game is exactly the PK-IND-game on the underlying KEM: $P_3 = (1 + \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}))/2$.

Hence, $\mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}'}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}}(\tau) + 4 \times \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\tau) + 4 \times \mathsf{Adv}^{\mathsf{ind}}_{\mathsf{PRG}_{\kappa, k+\ell}}(\tau)$. $\qquad\quad\square$

D.3 Proof of theorem 11

The proof is quite similar to above proof. In the selective setting, the adversary asks, from the beginning, the keys it wants to get, before seeing the global public parameters PK.

Game $\mathbf{G}_0$: In the initial game, the adversary thus asks for the keys it wants: for several sets A_j. One calls $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KEM}'.\mathsf{KeyGen}(1^\kappa)$, for each $S_i \in \Sigma$, and provides PK together with all the asked keys sk_i, for $i \in A = \cup A_j$ (all the asked sets). The adversary answers with two sets B_0 and B_1, but with the constraints that $A \cap B_0 = A \cap B_1 = \emptyset$ and $|B_0| = |B_1|$, and the challenger flips a random bit $b \xleftarrow{\$} \{0; 1\}$, generates a random session key $K \leftarrow \mathcal{K}$, runs $(C_i, K_i) \leftarrow \mathsf{KEM}'.\mathsf{Enc}(\mathsf{pk}_i)$ for all $i \in B_0 \cup B_1$, and outputs $C \leftarrow (C_i, E_i = K \oplus K_i)_{i \in B_b}$ together with the challenged key K. The adversary outputs its guess b'. We denote P_0 the probability of event $b' = b$, which is $(1 + \mathsf{Adv}^{\mathsf{ac\text{-}ind}}_{\mathsf{KEMAC}}(\mathcal{A}))/2$.

Game $\mathbf{G}_1$: In this game, one always outputs $C \leftarrow (C_i, E_i = K \oplus K_i)_{i \in B_0}$. To show this game is indistinguishable from the previous one, we define a sequence of hybrid games, for index I, such that for the I first $(C_i, E_i = K \oplus K_i)$ in C are for indices in B_0 and the last are for indices in B_b. For $I = 0$, this is $\mathbf{G}_0$, whereas for $I = |B_0| = |B_1|$, this is $\mathbf{G}_1$. And the gap between I and $I + 1$ is the

PK-IND-game on the underlying KEM'. Hence, $P_0 - P_1 \leq |B_0| \times \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}'}(\tau)$, where τ is the maximum running-time of adversary $\mathcal{A}$.

Furthermore, in this final game, this is clear that $P_1 = 1/2$, as the challenge C does not depend on b anymore.

Hence, $\mathsf{Adv}^{\mathsf{ac\text{-}ind}}_{\mathsf{KEMAC}}(\mathcal{A}) \leq 2S_B \times \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}'}(\tau)$, where S_B is the constant-size of the sets B. □

D.4 Proof of theorem 4

We present the sequence of games, first exploiting the session-key privacy of KEM_1.

Game $\mathbf{G_0}$: In the initial game, the adversary receives $\mathsf{pk} \leftarrow (\mathsf{pk}_1, \mathsf{pk}_2)$, both keys having been generated by the respective key generation algorithms, together with $C = (C_1, C_2)$ and K that is either $K_1 \oplus K_2$ (each encapsulated in C_1 and C_2) or a random key from $\mathcal{K}$, according to the random bit b. The adversary outputs its guess b'. We denote P_0 the probability of event $b' = b$, which is $(1 + \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}))/2$.

Game $\mathbf{G_1}$: In this game, we replace K_1's by $K_1 \xleftarrow{\$} \mathcal{K}$ in the generation of $K_1 \oplus K_2$: $P_0 - P_1 \leq \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}_1}(t)$, where t is the maximum running-time of adversary $\mathcal{A}$. In this final game, this is clear that $P_1 = 1/2$, as K is truly random in $\mathcal{K}$ in both cases.

Hence, $\mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}_1}(t)$. In the same way, one can prove $\mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}_2}(t)$, Hence $\mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}) \leq \min\{\mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}_1}(\tau), \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{KEM}_2}(\tau)\}$. □

D.5 Proof of theorem 5

The proof is quite similar to the previous one, but with the security of both schemes:

Game $\mathbf{G_0}$: In the initial game, the adversary receives $\mathsf{pk}^{(0)} \leftarrow (\mathsf{pk}_1^{(0)}, \mathsf{pk}_2^{(0)})$ and $\mathsf{pk}^{(1)} \leftarrow (\mathsf{pk}_1^{(1)}, \mathsf{pk}_2^{(1)})$, with keys having been generated by the respective key generation algorithms, together with $C = (C_1, C_2)$ and $K = K_1 \oplus K_2$ where $(C_i, K_i) \leftarrow \mathsf{KEM}_i.\mathsf{Enc}(\mathsf{pk}_i^{(b)})$, according to the random bit b. The adversary outputs its guess b'. We denote P_0 the probability of event $b' = b$, which is $(1 + \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}))/2$.

Game $\mathbf{G_1}$: In this game, we replace (C_1, K_1) by $(C_1, K_1) \leftarrow \mathsf{KEM}_1.\mathsf{Enc}(\mathsf{pk}_1^{(0)})$ in the generation of $C = (C_1, C_2)$ and $K = K_1 \oplus K_2$: $P_0 - P_1 \leq \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}_1}(t)$, where t is the maximum running-time of adversary $\mathcal{A}$.

Game $\mathbf{G_2}$: In this game, we replace (C_2, K_2) by $(C_2, K_2) \leftarrow \mathsf{KEM}_2.\mathsf{Enc}(\mathsf{pk}_2^{(0)})$ in the generation of $C = (C_1, C_2)$ and $K = K_1 \oplus K_2$: $P_1 - P_2 \leq \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}_2}(t)$. In this final game, this is clear that $P_2 = 1/2$, as (C, K) is independent of b.

Hence, we can claim $\mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}_1}(\tau) + \mathsf{Adv}^{\mathsf{pk\text{-}ind}}_{\mathsf{KEM}_2}(\tau)$. □

D.6 Proof of theorem 12

From a Diffie-Hellman tuple (A, B, C), one can derive, for random $s_k \xleftarrow{\$} \mathbb{Z}_q$, for $k = 1 \ldots, t+1$

$$ h_k \leftarrow g^{s_k} \qquad\qquad h \leftarrow B \qquad\qquad C_k \leftarrow A^{s_k} \qquad\qquad K \leftarrow C $$

where s, r are implicitly defined as $A = g^r$ and $B = g^s$. If C is indeed the Diffie-Hellman value for (g, A, B), then $K = C = g^{sr} = h^r$, we are in the real case ($b = 0$). If C is a random value, we are in the random case ($b = 1$):

$$ \mathsf{Adv}^{\mathsf{sk\text{-}ind}}_{\mathsf{TKEM}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{ddh}}_{\mathbb{G}}(\tau). $$

□

D.7 Proof of theorem 13

From a Diffie-Hellman tuple (A, B, C), one can derive, for random scalars $z_k^{(0)}, z_k^{(1)}, s_k^{(0)}, s_k^{(1)}, z^{(0)}, z^{(1)}, s^{(0)}, s^{(1)} \xleftarrow{\$} \mathbb{Z}_q$, for $k = 1 \dots, t+1$

$$h_k^{(0)} \leftarrow A^{z_k^{(0)}} \cdot g^{s_k^{(0)}} \qquad\qquad h^{(0)} \leftarrow A^{z^{(0)}} \cdot g^{s^{(0)}}$$
$$h_k^{(1)} \leftarrow A^{z_k^{(1)}} \cdot g^{s_k^{(1)}} \qquad\qquad h^{(1)} \leftarrow A^{z^{(1)}} \cdot g^{s^{(1)}}$$
$$C_k \leftarrow C^{z_k^{(b)}} \cdot B^{s_k^{(b)}}$$

where r is implicitly defined as $B = g^r$. If C is indeed the Diffie-Hellman value for (g, A, B), then $C_k = A^{rz_k^{(b)}} \cdot g^{rs_k^{(b)}} = (A^{z_k^{(b)}} \cdot g^{s_k^{(b)}})^r = (h_k^{(b)})^r$. If C is a random value $C = A^{r+c}$, for a random $c \xleftarrow{\$} \mathbb{Z}_q$:

$$C_k = A^{(r+c)z_k^{(b)}} \cdot g^{rs_k^{(b)}} = (A^{z_k^{(b)}} \cdot g^{s_k^{(b)}})^r \cdot A^{cz_k^{(b)}} = (h_k^{(b)})^r \cdot A^{cz_k^{(b)}}.$$

As $z_k^{(b)}$ is perfectly hidden in the public key, C_k follows a uniform distribution in $\mathbb{G}$, independently of the public key, and thus of b: $\mathsf{Adv}_{\mathsf{TKEM}}^{\mathsf{pk\text{-}ind}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathbb{G}}^{\mathsf{ddh}}(\tau)$. $\qquad\square$

D.8 Proof of theorem 17

In the selective setting, the adversary asks, from the beginning, the keys it wants to get, before seeing the global public parameters PK.

Game $\mathbf{G}_0$: In the initial game, the adversary thus asks for the keys it wants: for several sets A_j. One calls $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KEM.KeyGen}(1^\kappa)$, for each $S_i \in \Sigma$, and provides PK together with all the asked keys sk_i, for $i \in A = \cup A_j$ (all the asked sets). The adversary answers with a set B, but with the constraint that $A \cap B = \emptyset$, and the challenger flips a random bit $b \xleftarrow{\$} \{0; 1\}$, generates two random session keys $K_0', K_1' \leftarrow \mathcal{K}$, runs $(C_i, K_i) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk}_i)$ for all $i \in B$, and outputs $C \leftarrow (i, C_i, E_i = K_0' \oplus K_i)_{i \in B}$ together with the challenged key K_b' (that is either the really encapsulated key if $b = 0$ or a random key if $b = 1$). The adversary outputs its guess b'. We denote P_0 the probability of event $b' = b$, which is $(1 + \mathsf{Adv}_{\mathsf{KEMAC}}^{\mathsf{sk\text{-}ind}}(\mathcal{A}))/2$.

Game $\mathbf{G}_1$: In this game, we replace all the K_i's by $K_i \xleftarrow{\$} \mathcal{K}$ in the generation of E_i. To show this game is indistinguishable from the previous one, we define a sequence of hybrid games, for index I, such that for all $i < I$, one replaces K_i by a random element in $\mathcal{K}$. For $I = 1$, this is $\mathbf{G}_0$, whereas for $I = q_k + 1$, where q_k is the maximal number of keys, this is $\mathbf{G}_1$. And the gap between I and $I + 1$ is the $\mathsf{SK\text{-}IND}$-game on the underlying KEM. Hence, $P_0 - P_1 \leq q_k \times \mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{sk\text{-}ind}}(\tau)$, where τ is the maximum running-time of adversary $\mathcal{A}$.

Game $\mathbf{G}_2$: In this game, we replace all the E_i's by $E_i \xleftarrow{\$} \mathcal{K}$, which is perfectly indistinguishable from $K_0' \oplus K_i$ for a random K_i, under the group-law property. Hence, $P_1 = P_2$. In this final game, this is clear that $P_2 = 1/2$, as K_0' and K_1' do not appear anymore in C, and so K_b' is just a random key.

Hence, $\mathsf{Adv}_{\mathsf{KEMAC}}^{\mathsf{sk\text{-}ind}}(\mathcal{A}) \leq 2q_k \times \mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{sk\text{-}ind}}(\tau)$. $\qquad\square$

E Basic (Without Anonymity) Subset-Cover KEMAC

For efficiency considerations, we will focus on the subset-cover approach: during the Setup, one defines multiple sets S_i; when generating a user key usk_j, a list A_j of subsets if specified, which implicitly means user $U_j \in S_i$ for all $i \in A_j$; at encapsulation time, a target set T is given by B, such that $T = \cup_{i \in B} S_i$.

Intuitively, S_i's are subsets of the universe of users, and to specify the receivers, one encapsulates the key K for a covering of the target set T. A KEMAC, for a list Σ of sets S_i, can then be defined from any KEM in $\mathcal{K}$ that is a group with internal law denoted $\oplus$:

- $\mathsf{KEMAC.Setup}(\Sigma)$, for each $S_i \in \Sigma$, runs $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KEM.KeyGen}(1^\kappa)$. Then $\mathsf{PK} \leftarrow (\mathsf{pk}_i)_i$ and $\mathsf{MSK} \leftarrow (\mathsf{sk}_i)_i$;
- $\mathsf{KEMAC.KeyGen}(\mathsf{MSK}, A_j)$ defines the user's secret key $\mathsf{usk}_j \leftarrow (i, \mathsf{sk}_i)_{i \in A_j}$;
- $\mathsf{KEMAC.Enc}(\mathsf{PK}, B)$ generates a random session key $K \leftarrow \mathcal{K}$, and runs $(C_i, K_i) \leftarrow \mathsf{KEM.Enc}(\mathsf{pk}_i)$ for all $i \in B$, and outputs $C \leftarrow (i, C_i, E_i = K \oplus K_i)_{i \in B}$ together with the encapsulated key K;
- $\mathsf{KEMAC.Dec}(\mathsf{usk}_j, C)$ looks for $i \in \mathsf{usk}_j \cap C$, to run $K_i' \leftarrow \mathsf{KEM.Dec}(\mathsf{sk}_i, C_i)$ and output $K \leftarrow K_i' \oplus E_i$.

In terms of attributes, one can consider that each S_i is associated to an attribute a_i, and being in S_i for a user U_j means owning the attribute a_i. At encapsulation time, B lists the attributes that allow to decrypt: as soon as a_i is in B, any user U_j owning a_i can decrypt.

For the above scheme, we can claim the SK-IND security, but unfortunately not the AC-IND security. As attributes are known, the correctness of the KEM implies the correctness of the KEMAC.

Theorem 17 (Session-Key Privacy). *If the underlying KEM is SK-IND-secure, the above basic subset-cover KEMAC is SK-IND-secure, for selective key-queries:* $\mathsf{Adv}^{sk\text{-}ind}_{KEMAC}(\tau) \le 2q_k \times \mathsf{Adv}^{sk\text{-}ind}_{KEM}(\tau)$, *where* q_k *is the number of key-queries.*

As this proof uses a classical hybrid technique to replace each key by a random one, we defer the proof to the Appendix D.8. One notes that we need B to be provided in the ciphertext: indices i are given. We definitely exclude access-control privacy. In order to get anonymity, theses indices should not be given.

F ElGamal UTKEM constructions

F.1 ElGamal UTKEM in the Random Oracle Model

Here is the derived ElGamal TKEM in the ROM:

ROM.**UTKEM**.KeyGen($1^\kappa, n, t, g, \mathbb{G}, q$)**:** This algorithm returns the outputs (pk, $\{\mathsf{usk}_j\}_j$) of TKEM.KeyGen $(1^\kappa, n, t, g, \mathbb{G}, q)$;

ROM.**UTKEM**.Enc(pk)**:** this algorithm runs $(C, K) \leftarrow$ TKEM.Enc(pk), and then with $\mathcal{H} : \{0;1\}^* \to \{0;1\}^\mu$, a hash function modeled as a random oracle, it sets: $K' \leftarrow \mathcal{H}(K)$, and returns (C, K');

ROM.**UTKEM**.Dec(usk_j, C)**:** this algorithm runs $K \leftarrow$ TKEM.Dec(usk_j, C), and then returns $K' \leftarrow \mathcal{H}(K)$.

Proof that ROM.UTKEM is SK-IND and PK-IND-secure follows immediately from the proofs that TKEM is SK-IND and PK-IND-secure, in the ROM.

F.2 ElGamal UTKEM without the ROM

The Uniform-Keys TKEM without the ROM, in its derived authenticated KEM scheme (as it is derived in a non black-box way), is described with the following algorithms, with PRG : $\{0;1\}^\kappa \to \{0;1\}^{\ell+k}$ a PRG, and denoting $\mathbb{E} = \langle g \rangle$ and $\tilde{\mathbb{E}} = \langle \tilde{g} \rangle$ common twisted curves of respective prime orders $p, \tilde{p}$ over the finite field $\mathbb{F}_q$ such that q is on κ bits and $|q - 2^\kappa| \le 2^{\kappa/2}$, and absc the operation converting a curve element to the bitstring of its abscissa:

UTKEM'.KeyGen($1^\kappa, n, t$)**:** This algorithm runs $(\mathsf{pk}_0, \{\mathsf{usk}_{0,j}\}_j) \leftarrow$ TKEM.KeyGen($1^\kappa, n, t, g, \mathbb{E}, q$), $(\mathsf{pk}_1, \{\mathsf{usk}_{1,j}\}_j) \leftarrow$ TKEM.KeyGen($1^\kappa, n, t, \tilde{g}, \tilde{\mathbb{E}}, q$), and returns: $(\mathsf{pk} = (\mathsf{pk}_0, \mathsf{pk}_1), \{\mathsf{usk}_j = (\mathsf{usk}_{0,j}, \mathsf{usk}_{1,j})\}_j)$;

UTKEM'.Enc(pk)**:** this algorithm flips a coin $b \xleftarrow{\$} \{0;1\}$, runs: $(c_0, s_0) \leftarrow$ TKEM.Enc(pk_0), $(c_1, s_1) \leftarrow$ TKEM.Enc(pk_1), and: $K\|V \leftarrow$ PRG(absc(s_b)), and finally outputs: $(C = (c_0, c_1, V), K)$;

UTKEM'.Dec(usk_j, C)**:** this algorithm runs $s_0 \leftarrow$ TKEM.Dec($\mathsf{usk}_{0,j}, c_0$), $s_1 \leftarrow$ TKEM.Dec($\mathsf{usk}_{1,j}, c_1$), $K'_0\|V'_0 \leftarrow$ PRG(absc(s_0)), $K'_1\|V'_1 \leftarrow$ PRG(absc(s_1)), and checks whether $V = V'_0$, in which case it returns K'_0, else if $V = V'_1$, in which case it returns K'_1. Else, if none of these two conditions are met, the algorithm returns $\perp$.

One remarks that in UTKEM', correctness is not perfect, but ensured except with probability $2^{-1-\ell}$

Theorem 18 (Security of UTKEM'). *If TKEM is SK-IND and PK-IND-secure, and* PRG$\{0;1\}^\kappa \to \{0;1\}^{k+\ell}$ *is an IND-secure PRG, then UTKEM' is SK-IND, PK-IND and AUTH-secure.*

Proof. Let $n, t \in \mathbb{N}^*$.

SK-IND-security of UTKEM' We present a sequence of games, starting from the initial SK-IND-security game against UTKEM':

Game $\mathbf{G}_0$: In the initial game, setting: $(\mathsf{pk} = (\mathsf{pk}_0, \mathsf{pk}_1), \mathsf{sk} = \{\mathsf{usk}_j = (\mathsf{usk}_{0,j}, \mathsf{usk}_{1,j})\}_j) \leftarrow$ UTKEM'. KeyGen($1^\kappa, n, t$), the adversary is given pk; K_1 is drawn uniformly at random from $\{0;1\}^k$, as well as b from $\{0;1\}$, and $(C = (c_0, c_1, V), K_0) \leftarrow$ UTKEM'.Enc(pk), and the adversary is provided with C and K_b. The adversary outputs her guess b' for b. We denote P_0 the probability of event $b' = b$, which is $(1 + \mathsf{Adv}^{sk\text{-}ind}_{UTKEM'}(\mathcal{A}))/2$.

Game G_1: this game is identical to the previous one, except that now in the $\mathsf{UTKEM}'.\mathsf{Enc}(\mathsf{pk})$ operation, the s_0 generated with c_0 is replaced with a random element of $\mathbb{E}$, and the s_1 generated with c_1 is replaced with a random element of $\tilde{\mathbb{E}}$. The adversary then looses twice her advantage in the SK-IND security game of the underlying TKEM, and: $P_0 - P_1 \leq 2 \cdot \mathsf{Adv}_{\mathsf{TKEM}}^{\mathsf{sk\text{-}ind}}(t)$.

Game G_2: this game is identical to the previous one except that now, the $\mathsf{absc}(s_a)$ (with $a \in \{0;1\}$) used as an input to PRG to generate V is replaced with a random bit in $\{0;1\}^\kappa$. According to [CFGP06]'s corollary 11, since q is on κ bits and $|q - 2^\kappa| \leq 2^{\kappa/2}$, then the adversary looses an advantage bounded by $(1 + \sqrt{2}) \cdot 2^{-\kappa/2}$, and: $P_1 - P_2 \leq (1 + \sqrt{2}) \cdot 2^{-\kappa/2}$.

Game G_3: this game is the same as the previous one except that now, $K_0 \| V$ is replaced by $K_0 \| V \xleftarrow{\$} \{0;1\}^k$. The adversary then looses her advantage in the PRG IND-security game: $P_2 - P_3 \leq \mathsf{Adv}_{\mathsf{PRG}}^{\mathsf{ind}}(t)$. In this final game, this is clear that $P_3 = 1/2$, as K_b and V are then uniformly random.

Hence, $\mathsf{Adv}_{\mathsf{UTKEM}'}^{\mathsf{sk\text{-}ind}}(\mathcal{A}) \leq 4 \cdot \mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{sk\text{-}ind}}(t) + (1 + \sqrt{2}) \cdot 2^{1-\kappa/2} + 2 \cdot \mathsf{Adv}_{\mathsf{PRG}}^{\mathsf{ind}}(t)$.

PK-IND-security of UTKEM′ We present a sequence of games, starting from the initial PK-IND-security game against UTKEM':

Game G_0: In the initial game, setting:

$$(\mathsf{pk}_0 = (\mathsf{pk}_{0,0}, \mathsf{pk}_{0,1}), \mathsf{sk}_0 = \{\mathsf{usk}_{0,j}\}_j = (\{\mathsf{usk}_{0,0,j}\}_j, \{\mathsf{usk}_{0,1,j}\}_j)) \leftarrow \mathsf{UTKEM}'.\mathsf{KeyGen}(1^\kappa, n, t);$$

$$(\mathsf{pk}_1 = (\mathsf{pk}_{1,0}, \mathsf{pk}_{1,1}), \mathsf{sk}_1 = \{\mathsf{usk}_{1,j}\}_j = (\{\mathsf{usk}_{1,0,j}\}_j, \{\mathsf{usk}_{1,1,j}\}_j)) \leftarrow \mathsf{UTKEM}'.\mathsf{KeyGen}(1^\kappa, n, t),$$

and:

$$(C_0 = (c_{0,0}, c_{0,1}, V_0), K_0) \leftarrow \mathsf{UTKEM}'.\mathsf{Enc}(\mathsf{pk}_0);$$
$$(C_1 = (c_{1,0}, c_{1,1}, V_1), K_1) \leftarrow \mathsf{UTKEM}'.\mathsf{Enc}(\mathsf{pk}_1),$$

and drawing $b \xleftarrow{\$} \{0;1\}$, the adversary is provided with $\mathsf{pk}_0, \mathsf{pk}_1$, and C_b.
The adversary outputs her guess b' for b. We denote P_0 the probability of the event $b' = b$, which is $(1 + \mathsf{Adv}_{\mathsf{UTKEM}'}^{\mathsf{pk\text{-}ind}}(\mathcal{A}))/2$.

Game G_1: this game is identical to the previous one, except that now, in the $\mathsf{UTKEM}'.\mathsf{Enc}(\mathsf{pk}_b)$ operation to generate C_b, the $s_{b,0}$ and $s_{b,1}$ generated with $c_{b,0}$ and $c_{b,1}$ are replaced by respective group elements of $\mathbb{E}$ and $\tilde{\mathbb{E}}$. The adversary then looses twice her advantage in the SK-IND game of the underlying TKEM: $P_0 - P_1 \leq 2 \cdot \mathsf{Adv}_{\mathsf{TKEM}}^{\mathsf{sk\text{-}ind}}(t)$.

Game G_2: this game is identical to the previous one except that now, denoting $\mathsf{absc}(s_{b,a})$ the input to PRG used to generate V_b, $\mathsf{absc}(s_{b,a})$ is replaced with a uniformly random element of $\{0;1\}^\kappa$. According to [CFGP06]'s corollary 11, since q is on κ bits and $|q - 2^\kappa| \leq 2^{\kappa/2}$, then the adversary looses an advantage bounded by $(1 + \sqrt{2}) \cdot 2^{-\kappa/2}$, and: $P_1 - P_2 \leq (1 + \sqrt{2}) \cdot 2^{-\kappa/2}$.

Game G_3: this game is the same as the previous one, except that now, $K_b \| V_b$ is replaced by uniformly random $K_b \| V_b \xleftarrow{\$} \{0;1\}^{k+\ell}$. The adversary then looses her advantage in the PRG IND-security game: $P_2 - P_3 \leq \mathsf{Adv}_{\mathsf{PRG}}^{\mathsf{ind}}(t)$.

Game G_4: this game is identical to the previous one except that now one replaces: $(c_{b,0}, c_{b,1})$ in the adversary's input by $(c_{1-b,0}, c_{1-b,1})$. This game is perfectly indistinguishable from the previous one: $P_3 = P_4$, and the adversary's advantage in it is bounded by twice her advantage in the TKEM PK-IND security game, so: $P_4 \leq 2 \cdot \mathsf{Adv}_{\mathsf{TKEM}}^{\mathsf{pk\text{-}ind}}(t)$.

Hence, $\mathsf{Adv}_{\mathsf{UTKEM}'}^{\mathsf{pk\text{-}ind}}(\mathcal{A}) \leq 4 \cdot \mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{sk\text{-}ind}}(t) + 2 \cdot \mathsf{Adv}_{\mathsf{PRG}}^{\mathsf{ind}}(t) + 4 \cdot \mathsf{Adv}_{\mathsf{TKEM}}^{\mathsf{pk\text{-}ind}}(t) + (1 + \sqrt{2}) \cdot 2^{1-\kappa/2}$.

AUTH-security of UTKEM′ We present a sequence of games, starting from the initial AUTH-security game against UTKEM':

Game G_0: In the initial game, setting:

$$(\mathsf{pk}_0 = (\mathsf{pk}_{0,0}, \mathsf{pk}_{0,1}), \mathsf{sk}_0 = \{\mathsf{usk}_{0,j}\}_j = (\{\mathsf{usk}_{0,0,j}\}_j, \{\mathsf{usk}_{0,1,j}\}_j)) \leftarrow \mathsf{UTKEM}'.\mathsf{KeyGen}(1^\kappa, n, t);$$
$$(\mathsf{pk}_1 = (\mathsf{pk}_{1,0}, \mathsf{pk}_{1,1}), \mathsf{sk}_1 = \{\mathsf{usk}_{1,j}\}_j = (\{\mathsf{usk}_{1,0,j}\}_j, \{\mathsf{usk}_{1,1,j}\}_j)) \leftarrow \mathsf{UTKEM}'.\mathsf{KeyGen}(1^\kappa, n, t),$$

then: $(C = (c_0, c_1, V), K) \leftarrow \mathsf{UTKEM}'.\mathsf{Enc}(\mathsf{pk}_0)$, one then runs $K' \leftarrow \mathsf{UTKEM}'.\mathsf{Dec}(\mathsf{sk}_1, C)$. We denote P_0 the probability that $K' \neq \perp$, which is $\mathsf{Adv}_{\mathsf{UTKEM}'}^{\mathsf{auth}}(\mathcal{A})$.

Game G_1: this game is identical to the previous one except that now in the $\mathsf{UTKEM'.Enc(pk_0)}$ operation providing C, the s_0 and s_1 are respectively replaced by uniformly random elements of $\mathbb{E}$ and $\tilde{\mathbb{E}}$. The difference between this game and the previous one is twice the $\mathsf{SK\text{-}IND}$ security game of TKEM, against a trivial adversary $\mathcal{A}_0$, so: $P_0 - P_1 \leq 2 \cdot \mathsf{Adv}_{\mathsf{TKEM'}}^{\mathsf{sk\text{-}ind}}(\tau)$, where τ is the running time of $\mathcal{A}_0$ computing two key generations, one encapsulation, and one decapsulation of $\mathsf{UTKEM'}$.

Game G_2: this game is identical to the previous one, except that now, the $\mathsf{absc}(s_b)$ used as input to PRG to generate V is replaced by a uniformly random element of $\{0;1\}^\kappa$. According to [CFGP06]'s corollary 11, since q is on κ bits and $|q - 2^\kappa| \leq 2^{\kappa/2}$, $P_1 - P_2 \leq (1 + \sqrt{2}) \cdot 2^{-\kappa/2}$.

Game G_3: this game is identical to the previous one, except that now one replaces the original $K\|V$ in the $\mathsf{UTKEM'.Enc(pk_0)}$ operation by $K\|V \xleftarrow{\$} \{0;1\}^{k+\ell}$. This is indistinguishable from the previous game except with probability $\mathsf{Adv}_{\mathsf{PRG}}^{\mathsf{ind}}(\tau')$, where τ' is the running time of another trivial adversary $\mathcal{A}_1$, running two key generations, one encapsulation and one decapsulation of $\mathsf{UTKEM'}$, so: $P_2 - P_3 \leq \mathsf{Adv}_{\mathsf{PRG}}^{\mathsf{ind}}(\tau')$. In this game, the probability that in the $\mathsf{UTKEM'.Dec(sk_1, C)}$ operation, the V_0' drawn from c_0 is equal to V is of $2^{-\ell}$, and the probability that the V_1' drawn from c_1 is equal to V is of $2^{-\ell}$ too, so the probability that $V_0' = V$ or $V_1' = V$, and so that something different from $\bot$ is returned is bounded by $2^{1-\ell}$, so: $P_3 \leq 2^{1-\ell}$.

Hence, $\mathsf{Adv}_{\mathsf{UTKEM'}}^{\mathsf{auth}}(\mathcal{A}) \leq 2^{1-\ell} + 4 \cdot \mathsf{Adv}_{\mathsf{KEM}}^{\mathsf{sk\text{-}ind}}(t) + (1 + \sqrt{2}) \cdot 2^{1-\kappa/2} + 2 \cdot \mathsf{Adv}_{\mathsf{PRG}}^{\mathsf{ind}}(t)$. $\qquad\square$

G   CRYSTALS-Kyber

G.1   Correctness

We can first show that decryption indeed gives back K. To this aim, one can note that we have

$$
\begin{aligned}
w = v - \mathbf{s}^T \mathbf{u} &= (\mathbf{b}^T \mathbf{r} + e_2 + \lceil \tfrac{q}{2} \rfloor \cdot K) - \mathbf{s}^T \cdot (\mathbf{A}^T \mathbf{r} + \mathbf{e}_1) \\
&= (\mathbf{s}^T \mathbf{A}^T + \mathbf{e}^T) \cdot \mathbf{r} + e_2 + \lceil \tfrac{q}{2} \rfloor \cdot K - \mathbf{s}^T \cdot (\mathbf{A}^T \mathbf{r} + \mathbf{e}_1) \\
&= \mathbf{s}^T \mathbf{A}^T \mathbf{r} + \mathbf{e}^T \mathbf{r} + e_2 + \lceil \tfrac{q}{2} \rfloor \cdot K - \mathbf{s}^T \mathbf{A}^T \mathbf{r} + \mathbf{s}^T \mathbf{e}_1 \\
&= \mathbf{e}^T \mathbf{r} + e_2 + \lceil \tfrac{q}{2} \rfloor \cdot K + \mathbf{s}^T \mathbf{e}_1 = \lceil \tfrac{q}{2} \rfloor \cdot K + z
\end{aligned}
$$

with $z = \mathbf{e}^T \mathbf{r} + \mathbf{s}^T \mathbf{e}_1 + e_2$. Then,

$$
\lceil \tfrac{2}{q} \cdot w \rfloor = \lceil \tfrac{2}{q} \cdot \left(\lceil \tfrac{q}{2} \rfloor \cdot K + z \right) \rfloor .
$$

As $K \in \{0;1\}^n$, if $\|z\|_\infty < q/4$, then $\lceil \tfrac{2}{q} \cdot w \rfloor = K$. And since $\mathbf{e}, \mathbf{r}, \mathbf{s}$ are drawn according to $\mathcal{B}_{\eta_1}$, and $\mathbf{e}_1, e_2$ according to $\mathcal{B}_{\eta_2}$, with appropriate n, k, q, η_1, η_2, $\|z\|_\infty < q/4$, with overwhelming probability [ABD+21], even with the additional compression function, that introduces an additional noise. And we omit it in the current description, as it is for efficiency purpose but does not impact the security notions we are interested in.

G.2   Indistinguishability of Kyber

We proceed with a sequence of games, so that the final game does not leak any information about the bit b used by the challenger when encrypting K_b among K_0 and K_1 chosen by the adversary.

Game G_0: The initial game is the security experiment, where a public key $\mathsf{pk} = (\mathbf{A}, \mathbf{b})$ is generated, as well as a ciphertext C of K_b.

Game G_1: We change the way the challenger generates the public key pk, with randomly chosen $\mathbf{b} \xleftarrow{\$} \mathsf{R}_q$. The distance of the distributions of the outputs of the adversary is then bounded by $\mathsf{Adv}_{\mathsf{R}_q, k, k, \eta_1}^{\mathsf{dmlwe}}(\tau)$. But still, the ciphertext is built as

$$
\begin{bmatrix} \mathbf{u} \\ v \end{bmatrix} = \left[\mathbf{A} | \mathbf{b} \right]^T \times \mathbf{r} + \begin{bmatrix} \mathbf{e}_1 \\ e_2 \end{bmatrix} + \begin{bmatrix} 0 \\ \lceil \tfrac{q}{2} \rfloor \cdot K_b \end{bmatrix} .
$$

Game G_2: We change now the simulation of the ciphertext, with $(\mathbf{w}, z) \xleftarrow{\$} \mathsf{R}_q^k \times \mathsf{R}_q = \mathsf{R}_q^{k+1}$, and

$$\begin{bmatrix} \mathbf{u} \\ v \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{w} \\ z \end{bmatrix} + \begin{bmatrix} 0 \\ \lceil \frac{q}{2} \rceil \cdot K_b \end{bmatrix}.$$

The distance of the distributions of the outputs of the adversary is then bounded by $\mathsf{Adv}_{\mathsf{R}_q, k+1, k, \eta_2}^{\mathsf{dmlwe}}(\tau)$, as $\mathbf{r} \xleftarrow{\$} \mathcal{B}_{\eta_1}^k$ and $(\mathbf{e}_1, e_2) \xleftarrow{\$} \mathcal{B}_{\eta_2}^k \times \mathcal{B}_{\eta_2} = \mathcal{B}_{\eta_2}^{k+1}$, and $\eta_1 \geq \eta_2$.

Game G_3: We eventually replace C by a random sampling in $\mathsf{R}_q^k \times \mathsf{R}_q$, which is perfectly indistinguishable. And then, the ciphertext is perfectly independent from K_b.

H   Some Figures

KEM$'$.KeyGen(1^κ):
1. $(\mathsf{pk}, \mathsf{sk}) \leftarrow$ KEM.KeyGen(1^κ)
2. return $(\mathsf{pk}, \mathsf{sk})$

KEM$'$.Enc(pk):
1. $(c, s) \leftarrow$ KEM.Enc(pk)
2. $U \| V \leftarrow \mathcal{H}(s)$
3. $C \leftarrow (c, V)$, $K \leftarrow U$
4. return (C, K)

KEM$'$.Dec($\mathsf{sk}, \mathsf{pk}, C = (c, V)$):
1. $s \leftarrow$ KEM.Dec(sk, c)
2. $U' \| V' \leftarrow \mathcal{H}(s)$
3. if $V = V'$, return U', otherwise reject

Fig. 1. Key-Confirmation KEM$'$

KEM.KeyGen(1^κ):
1. $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow$ KEM.KeyGen$_i$(1^κ), for $i = 1, 2$
2. $\mathsf{pk} \leftarrow (\mathsf{pk}_1, \mathsf{pk}_2)$; $\mathsf{sk} \leftarrow (\mathsf{sk}_1, \mathsf{sk}_2)$
3. return $(\mathsf{pk}, \mathsf{sk})$

KEM.Enc(pk):
1. $(C_i, K_i) \leftarrow$ KEM$_i$.Enc(pk_i), for $i = 1, 2$
2. $C \leftarrow (C_1, C_2)$; $K \leftarrow K_1 \oplus K_2$
3. return (C, K)

KEM.Dec(sk, C):
1. $K_i \leftarrow$ KEM$_i$.Dec(sk_i, C_i), for $i = 1, 2$
2. $K \leftarrow K_1 \oplus K_2$
3. return K

Fig. 2. Hybrid KEM

H.1   Anonymity of Kyber

We will show that the ciphertexts are statistically close to uniform in the ciphertext space, i.e. in $\mathsf{R}_q^k \times \mathsf{R}_q$, whatever the public key.

Game G_0: The initial game is the security experiment, where two public keys $\mathsf{pk}_0 = (\mathbf{A}_0, \mathbf{b}_0)$ and $\mathsf{pk}_1 = (\mathbf{A}_1, \mathbf{b}_1)$ are generated and a ciphertext C is generated according to pk_b, an a message K chosen by the adversary.

Game G_1: We change the way the challenger generates the public keys pk_0 and pk_1, with randomly chosen $\mathbf{b}_0, \mathbf{b}_1 \overset{\$}{\leftarrow} \mathsf{R}_q$. The distance of the distributions of the outputs of the adversary is then bounded by $2 \times \mathsf{Adv}^{\mathsf{dmlwe}}_{\mathsf{R}_q, k, k, \eta_1}(\tau)$. But still, the ciphertext is built as

$$\begin{bmatrix} \mathbf{u} \\ v \end{bmatrix} = \begin{bmatrix} \mathbf{A}_b | \mathbf{b}_b \end{bmatrix}^T \times \mathbf{r} + \begin{bmatrix} \mathbf{e}_1 \\ e_2 \end{bmatrix} + \begin{bmatrix} 0 \\ \lceil \frac{q}{2} \rfloor \cdot K \end{bmatrix}.$$

Game G_2: We change now the simulation of the ciphertext, with $(\mathbf{w}, z) \overset{\$}{\leftarrow} \mathsf{R}_q^k \times \mathsf{R}_q = \mathsf{R}_q^{k+1}$, and

$$\begin{bmatrix} \mathbf{u} \\ v \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{w} \\ z \end{bmatrix} + \begin{bmatrix} 0 \\ \lceil \frac{q}{2} \rfloor \cdot K \end{bmatrix}.$$

The distance of the distributions of the outputs of the adversary is then bounded by $\mathsf{Adv}^{\mathsf{dmlwe}}_{\mathsf{R}_q, k+1, k, \eta_2}(\tau)$, as $\mathbf{r} \overset{\$}{\leftarrow} \mathcal{B}^k_{\eta_1}$ and $(\mathbf{e}_1, e_2) \overset{\$}{\leftarrow} \mathcal{B}^k_{\eta_2} \times \mathcal{B}_{\eta_2} = \mathcal{B}^{k+1}_{\eta_2}$, and $\eta_1 \geq \eta_2$.

Game G_3: We eventually replace C by a random sampling in $\mathsf{R}_q^k \times \mathsf{R}_q$, which is perfectly indistinguishable. And then, the ciphertext is perfectly independent from the used public key.

H.2 Parameters

In Table 3, we recall the parameters of Kyber, where δ is the security level (one can see $\delta = 2^{-\kappa}$), and (d_u, d_v) are some parameters for the compression function, with the sizes (in Bytes) of the keys and ciphertexts, using the compression/decompression.

Table 3. Parameter Sets for Kyber

Parameters	n	k	q	η_1	η_2	(d_u, d_v)	pk	sk	c	δ
Kyber-512	256	2	3329	3	2	(10,4)	800B	32B	736B	2^{-139}
Kyber-768	256	3	3329	2	2	(10,4)	1184B	32B	1088B	2^{-164}
Kyber-1024	256	4	3329	2	2	(11,5)	1568B	32B	1568B	2^{-174}