

On the Complete Non-Malleability of the Fujisaki-Okamoto Transform*

Daniele Friolo^{†1}, Matteo Salvino^{‡2}, and Daniele Venturi^{§1}

¹Department of Computer Science, Sapienza University of Rome, Rome, Italy

²Research Institute CODE, Universität der Bundeswehr München, Munich, Germany

Abstract

The Fujisaki-Okamoto (FO) transform (CRYPTO 1999 and JoC 2013) turns any weakly (i.e., IND-CPA) secure public-key encryption (PKE) scheme into a strongly (i.e., IND-CCA) secure key encapsulation method (KEM) in the random oracle model (ROM). Recently, the FO transform re-gained momentum as part of CRISTAL-Kyber, selected by the NIST as the PKE winner of the post-quantum cryptography standardization project.

Following Fischlin (ICALP 2005), we study the *complete non-malleability* of KEMs obtained via the FO transform. Intuitively, a KEM is completely non-malleable if no adversary can maul a given public key and ciphertext into a new public key and ciphertext encapsulating a related key for the underlying blockcipher.

On the negative side, we find that KEMs derived via FO are *not* completely non-malleable in general. On the positive side, we show that complete non-malleability holds in the ROM by assuming the underlying PKE scheme meets an additional property, or by a slight tweak of the transformation.

*This is the full version of an accepted paper of ACNS '23. The work of the first and third author is supported by the SPECTRA Grant of Sapienza University of Rome.

[†]friolo@di.uniroma1.it

[‡]matteo.salvino@unibw.de —The work was carried out whilst the author was a student at Sapienza University of Rome, Rome, IT.

[§]venturi@di.uniroma1.it

Contents

1	Introduction	3
1.1	Our Contributions	3
1.2	Related Work	5
2	Preliminaries	5
2.1	Notation	5
2.2	Public-Key Encryption	6
2.2.1	OW-PCA and OW-PCVA	6
2.2.2	Complete Non-Malleability	7
2.3	Secret-Key Encryption	7
3	Completely Non-Malleable KEMs	8
3.1	Relation between NM-ATK* and SNM-ATK*	10
4	Analysis of Fujisaki-Okamoto transforms	13
4.1	Analysis of the $U_m^{\perp/\neq}$ transformations	14
4.2	Analysis of the $U^{\perp/\neq}$ transformations	14
4.3	Modified transformation \hat{U}^{\perp}	19
5	Relation with Completely Non-Malleable PKE	22
5.1	NM-ATK* PKE \implies NM-ATK* KEM	22
5.2	NM-ATK* KEM + NM-ATK SKE \implies NM-ATK* PKE	23

1 Introduction

Public-key encryption (PKE) allows Alice to encrypt a message under a Bob’s public key, so that Bob can decrypt the ciphertext using the corresponding secret key. Several security notions for PKE have been proposed in the literature. The most basic one, namely *indistinguishability against chosen-plaintext attacks* (IND-CPA) requires that an adversary, given the public key, cannot distinguish between the encryption of two messages.

Non-malleability. As noted for the first time by Dolev *et al.* [9], IND-CPA appears to be insufficient for many applications. Consider for instance the setting of private auctions. Here, a bidder can sample its own pair of public/secret keys, encrypt the bid b using the public key, and send the encryption together with the public key to the auctioneer. After all the participants have sent their bid, the auctioneer can declare the winner by asking each party to reveal the secret key (or the bid itself, along with the random coins used for encryption). A malicious user, given a ciphertext c containing the bid of another party, can try to construct a ciphertext c' that, when decrypted, leads to a bid b' such that $b' > b$.

In light of such malleability attacks, stronger security notions for PKE schemes have been introduced. These include the notions of *non-malleability under chosen-plaintext* and *chosen-ciphertext attacks* [9, 3, 6, 18] (NM-CPA and NM-CCA), and *indistinguishability under chosen-ciphertext attacks* (IND-CCA). All of these notions imply that the attacker, given the public key and a target ciphertext, is unable to craft a mauling ciphertext whose underlying plaintext is related to the one contained in the target ciphertext.

Complete non-malleability. In 2005, Fischlin [12] noted that non-malleability might be still insufficient for some applications. In fact, the above notions do not account for the possibility that the attacker may try to maul the public key as well. For instance, consider again the setting of private auctions. A malicious user, knowing a ciphertext c and the public key pk , may try to craft a public key pk' and a ciphertext c' which encrypt a bid $b' > b$. To capture these attacks, Fischlin introduced *complete non-malleability*, which rules out such adversaries.

As noted by Fischlin himself, completely non-malleable PKE has several useful applications, including key-agreement protocols with security against unknown key attacks, and signature schemes with security against strong unforgeability attacks (as needed, e.g., in e-cash systems).

Known constructions. Fischlin [12] showed that a simple variant of RSA-OAEP is completely non-malleable in the random oracle model (ROM).¹ Ventre and Visconti [22] later gave two constructions of completely non-malleable PKE in the common reference string (CRS) model, based on non-interactive zero-knowledge (NIZK) proofs for all of NP.

Subsequent work provided more efficient constructions of completely non-malleable PKE without random oracles, using both pairing-based assumptions [7, 15] and lattice-based assumptions [20, 21].

1.1 Our Contributions

In practice, due to its computational overhead, PKE is never used to encrypt long messages. Rather, as it happens in many real-world protocols (including TLS), the parties use public-key techniques in order to establish a common secret key for a blockcipher, which can be later used in order to encrypt any subsequent communication of arbitrary length. In the literature this paradigm is also known as the *key/data encapsulation method* (KEM/DEM), or simply hybrid encryption. This motivates our main question:

Can we get efficient constructions of completely non-malleable PKE via the KEM/DEM paradigm?

¹He also proves that the original version of RSA-OAEP, as well as the Cramer-Shoup PKE [8], is not completely non-malleable.

Our main contribution is a positive answer to the above question. Namely, we put forward natural notions of complete non-malleability for KEMs and show that these notions are sufficient to imply completely non-malleable PKE with small ciphertext rate. Furthermore, we show that an already existing, widely-used, KEM meets our notions. We elaborate on these contributions below.

Definitions. A *key encapsulation method* (KEM) is made of two algorithms: An encapsulation algorithm that, given the public key pk , outputs a ciphertext c encapsulating a secret key K ; and a decapsulation algorithm that, given the secret key sk corresponding to pk , allows to recover K . Similarly to PKE, several non-malleability properties for KEMs have been introduced.

In [Section 3](#), we put forward three indistinguishability-based variants of completely non-malleable KEMs (dubbed NM-CPA*, NM-CCA1* and NM-CCA2*), capturing different flavors of chosen-plaintext and chosen-ciphertext attacks. In [Section 3.1](#), we define the corresponding simulation-based variants (dubbed SNM-CPA*, SNM-CCA1* and SNM-CCA2*), and show the equivalence between the NM-ATK* and SNM-ATK* notions for $\text{ATK} \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$. More specifically, we show that for NM-CPA* and SNM-CPA* the equivalence holds for so-called *complete relations*, while for NM-CCA1* and SNM-CCA1*, and for NM-CCA2* and SNM-CCA2*, the equivalence holds for a restricted set of relations called *lacking relations* (see [Section 3.1](#) for details). These findings are in line with the work by Ventre and Visconti [[22](#)], who showed analogous results for completely non-malleable PKE.

Analysis of Fujiasaki-Okamoto. As our main contribution, we analyze the complete non-malleability of the Fujiasaki-Okamoto (FO) transform [[13](#)]. Recall that the FO transform turns any IND-CPA secure PKE into an IND-CCA secure KEM in the ROM, without affecting the ciphertext size, and at the cost of a very small extra computation effort w.r.t. the underlying PKE scheme (when the RO is replaced with a real-world hash function like SHA-256). Recently, the FO transform re-gained momentum as part of CRISTAL-Kyber, selected by the NIST as the PKE winner of the post-quantum cryptography standardization project [[17](#)]. In this light, we believe that investigating further security properties of the FO transform is a very natural research question.

Our analysis follows the modular analysis of the FO transform due to Hofheinz *et al.* [[14](#)]. Here, one interprets the FO transform as a sequence of two transformations T and U :

- The transformation T starts with any IND-CPA PKE. The encryption algorithm runs the encryption algorithm of the underlying PKE scheme but sets its randomness to $G(m)$, where G is a RO. The decryption algorithm runs the decryption algorithm of the underlying PKE scheme, and returns \perp if the decrypted message m' is \perp or if the encryption of m' with randomness $G(m')$ does not equal the ciphertext.
- The transformation U takes a PKE scheme satisfying different flavours of one-wayness (which are achieved by the transformation T), and outputs an IND-CCA secure KEM. This transformation essentially comes in 2 variants.² U_m calculates the encapsulated key K by randomly choosing a message m from the message space of the underlying PKE scheme, encrypting m under pk , and then computing K as $H(m)$, where H is a RO. U instead computes the key as $H(m, c)$.

First, in [Section 4.1](#), we show a concrete attack against the transformation U_m that works even when considering the weakest flavour of complete non-malleability (i.e., NM-CPA*). We take the El-Gamal PKE scheme as the base PKE scheme to be transformed by T and U into a KEM scheme. In particular, we prove that an adversary can appropriately maul the public key pk and the ciphertext c encrypting m , and come up with a ciphertext c' encrypting the same message m under a different public key pk' . Since the encapsulated key computed by U_m is $H(m)$, the key encapsulated by c and c' will be the same. Thus, complete non-malleability is trivially broken.

Second, in [Section 4.2](#), we show that the transformation U is not completely non-malleable. To see this, it suffices to take a contrived PKE scheme in which we add a dummy bit to the public key of a PKE scheme

²Each of U and U_m also comes in 2 variants, but the difference is irrelevant for what follows.

satisfying the one-wayness properties required by the transformation U . This additional bit is completely ignored by the encryption algorithm and does not effect one-wayness. However, one can trivially break complete non-malleability by flipping the last bit of the public key. On the positive side, we show that U does achieve complete non-malleability assuming the underlying PKE scheme meets a natural public-key uniqueness property, where the latter essentially means that an adversary cannot come up with different public keys $\mathsf{pk}, \mathsf{pk}'$ for which there exist a message m and a ciphertext c such that c is a valid encryption of m under both pk and pk' .

Indeed, we point out that uniqueness seems to be a standard property to achieve non-malleability, e.g. quasi-unique responses for Fiat-Shamir signatures [12] and [11].

Finally, in Section 4.3, we show how to tweak the transform U in order to obtain complete non-malleability without requiring public-key uniqueness. For this, it suffices to compute the key K as $\mathsf{H}(m, c, \mathsf{pk})$. This way, even if the attacker can break public-key uniqueness, the random oracle will ensure that the two encapsulated keys are independent. We notice that a similar technique was already used in [10], where regular CCA security of the FO transform in the multi-user setting is achieved by adding just a fraction of the public key with high min-entropy as an input to the hash function. However, it is not clear whether complete non-malleability of FO is achievable with such a slight modification.

Relation with completely non-malleable PKE. In Section 5, we show that by combining a completely non-malleable KEM with a non-malleable secret-key encryption (SKE) scheme we obtain a completely non-malleable PKE using the KEM/DEM paradigm. Furthermore, we observe that one can always obtain a completely non-malleable KEM by encrypting a random secret key via a completely non-malleable PKE.

1.2 Related Work

Nagao *et al.* in [16] analyze standard non-malleability in the context of key encapsulation. In particular, they consider different flavours of non-malleable KEM, such as NM-CPA, NM-CCA1 and NM-CCA2.

Ventre and Visconti [22] note that the stronger CCA2 notion of complete non-malleability is not strictly necessary for some of the applications proposed by Fischlin [12]. Hence, they put forward weaker flavours of complete non-malleability, and establish the relations between the standard comparison-based notions NM-ATK* and their simulation-based counterparts SNM-ATK* for $\text{ATK} \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$. They also give two constructions of completely non-malleable NM-CCA2* secure PKE: one in the CRS model using NIZK proofs for all of NP, and one in the plain model using interactive encryption.

Barbosa and Farshim [1] consider an equivalent indistinguishability-based notion of complete non-malleability based on so-called strong CCA security, in which the (strong) decryption oracle provides decryptions under arbitrarily chosen public keys. Duman *et al.* [10] analyze CCA security of the FO transform in the multi-user setting.

2 Preliminaries

In this section we introduce some basic notation and recall a few standard definitions that will be used later to prove some of our results.

2.1 Notation

We use calligraphic letters to denote sets, such as \mathcal{X} , and lower-case letters for variables, such as x . We use $x \leftarrow_s \mathcal{X}$ to indicate that x is picked uniformly at random from \mathcal{X} . A similar notation is used in the presence of a randomized or probabilistic algorithm A . Indeed, $x \leftarrow_s \mathsf{A}(\cdot)$ means that x is the output of the randomized algorithm A . Alternatively, when a random coin r is given as an input of A , we equivalently write $x := \mathsf{A}(\cdot; r)$. All the algorithms we will consider are PPT (Probabilistic Polynomial Time), i.e. for any input $x \in \{0, 1\}^*$ and random coin r , $\mathsf{A}(x; r)$ terminates in at most polynomial many steps, in the size of its inputs. When an algorithm A has access to a set of oracles $\{\mathsf{O}_1, \dots, \mathsf{O}_n\}$, we use the notation $\mathsf{A}^{\mathsf{O}_1, \dots, \mathsf{O}_n}$, to indicate

that A can interact in a black-box manner with oracles O_1, \dots, O_n during its computation. We denote with $\lambda \in \mathbb{N}$ the security parameter and we will assume that all the algorithms we will consider take λ as an input. A function $\nu : \mathbb{N} \rightarrow [0, 1]$ is negligible if for every polynomial $p(n) \exists N \in \mathbb{N}$ s.t. $\forall n_0 \geq N, \nu(n_0) < \frac{1}{p(n_0)}$. We denote with $\text{negl}(\lambda)$ any function that is negligible in λ . Given two random variables X and Y , we denote $X \approx_c Y$ when X and Y are computationally indistinguishable, and with $X \equiv Y$ when X and Y are identically distributed.

2.2 Public-Key Encryption

A public-key encryption (PKE) scheme Π consists of three algorithms ($\text{Gen}, \text{Enc}, \text{Dec}$), together with a message space \mathcal{M} (which we assume to be efficiently recognizable) where:

- The key generation algorithm Gen takes as input 1^λ and outputs a public-private key pair (pk, sk) .
- The encryption algorithm Enc takes as inputs a public key pk and a message $m \in \mathcal{M}$, and outputs an encryption c of the message m under pk .
- The deterministic decryption algorithm Dec takes as inputs a decryption key sk and a ciphertext c , and outputs either a message $m \in \mathcal{M}$, or \perp (denoting failure).

Next, we define both correctness and security of PKE as needed for our purposes. Some of the definitions below are taken verbatim from [14].

Definition 1 (γ -uniformity). Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be a PKE scheme with message space \mathcal{M} . Given $(\text{pk}, \text{sk}) \leftarrow_s \text{Gen}(1^\lambda)$, a message $m \in \mathcal{M}$ and a ciphertext c we define the γ -uniformity function as follows

$$\gamma(m, c) = \Pr [c = \text{Enc}(\text{pk}, m)],$$

where the probability is taken over the choice of the random coins used for encrypting m under pk .

We say that Π is γ -uniform if, for any $(\text{pk}, \text{sk}) \in \text{Gen}(1^\lambda)$, any message $m \in \mathcal{M}$, and any ciphertext $c \in \{0, 1\}^*$, it holds that $\gamma(m, c) \leq \gamma$.

Definition 2 (correctness). Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be a PKE scheme with message space \mathcal{M} . A PKE scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is correct if $\forall m \in \mathcal{M}, \forall (\text{pk}, \text{sk}) \leftarrow_s \text{Gen}(1^\lambda)$

$$\Pr [\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m)) = m] = 1.$$

2.2.1 OW-PCA and OW-PCVA

We recall the definitions of one-wayness under plaintext checking attacks (OW-PCA) and one-wayness under plaintext and validity checking attacks (OW-PCVA).

Definition 3 (OW-ATK). Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be a public-key encryption scheme with message space \mathcal{M} . We define the following $\text{PKE}^{\text{ow-atk}}$ games for $\text{atk} \in \{\text{pca}, \text{pcva}\}$

<p>Experiment $\text{PKE}_{\Pi, A}^{\text{ow-atk}}(\lambda)$</p> <p>$(\text{pk}^*, \text{sk}^*) \leftarrow_s \text{Gen}(1^\lambda)$</p> <p>$m^* \leftarrow_s \mathcal{M}$</p> <p>$c^* \leftarrow_s \text{Enc}(\text{pk}^*, m^*)$</p> <p>$m' \leftarrow_s A^{O_1}(\text{pk}, c^*)$</p> <p>return $\text{PCO}(\text{sk}^*, m', c^*)$</p>	<p>Oracle $\text{PCO}(\text{sk}^*, m, c)$</p> <hr style="border: 0.5px solid black;"/> <p>return 1 iff</p> <p>$(\text{Dec}(\text{sk}^*, c) = m) \wedge (m \neq \perp)$</p> <p>Oracle $\text{CVO}^{(c^*)}(\text{sk}^*, c)$</p> <hr style="border: 0.5px solid black;"/> <p>$m := \text{Dec}(\text{sk}^*, c)$</p> <p>return 1 iff $m \neq \perp$</p>
--	---

In the experiment above,

$$\begin{aligned} \text{if } \text{atk} = \text{pca} \text{ then } \mathcal{O}_1 &= \text{PCO}(\text{sk}^*, \cdot, \cdot), \\ \text{if } \text{atk} = \text{pcva} \text{ then } \mathcal{O}_1 &= \text{PCO}(\text{sk}^*, \cdot, \cdot), \text{CVO}^{(c^*)}(\text{sk}^*; \cdot), \end{aligned}$$

where $\text{CVO}^{(c^*)}(\text{sk}, \cdot)$ means that \mathcal{A} is allowed to query the CVO algorithm for any ciphertext distinct from the challenge ciphertext c^* . We say that Π is OW-ATK secure if for all PPT \mathcal{A} ,

$$\Pr \left[\text{PKE}_{\Pi, \mathcal{A}}^{\text{ow-atk}}(\lambda) = 1 \right] \leq \text{negl}(\lambda).$$

2.2.2 Complete Non-Malleability

Finally, we recall the indistinguishability-based security definition for completely non-malleable PKE as defined by Ventre and Visconti [22].

Definition 4 (NM-CPA*, NM-CCA1*, NM-CCA2*). Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be a public-key encryption scheme, and let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a PPT adversary. For $\text{atk} \in \{\text{cpa}, \text{cca1}, \text{cca2}\}$ and $\lambda \in \mathbb{N}$ let

$$\text{PKE}_{\Pi, \mathcal{A}}^{\text{nm-atk}^*}(\lambda) \approx_c \text{PKE}_{\Pi, \mathcal{A}, \$}^{\text{nm-atk}^*}(\lambda),$$

where the experiments $\text{PKE}_{\Pi, \mathcal{A}}^{\text{nm-atk}^*}(\lambda)$ and $\text{PKE}_{\Pi, \mathcal{A}, \$}^{\text{nm-atk}^*}(\lambda)$ are defined as follows:

Experiment $\text{PKE}_{\Pi, \mathcal{A}}^{\text{nm-atk}^*}(\lambda)$	Experiment $\text{PKE}_{\Pi, \mathcal{A}, \$}^{\text{nm-atk}^*}(\lambda)$
$(\text{pk}^*, \text{sk}^*) \leftarrow_s \text{Gen}(1^\lambda)$	$(\text{pk}^*, \text{sk}^*) \leftarrow_s \text{Gen}(1^\lambda)$
$(\mathcal{M}, s) \leftarrow_s \mathcal{A}_1^{\text{O}_1}(\text{pk})$	$(\mathcal{M}, s) \leftarrow_s \mathcal{A}_1^{\text{O}_1}(\text{pk})$
$m^* \leftarrow_s \mathcal{M}$	$m^*, \tilde{m} \leftarrow_s \mathcal{M}$
$c^* \leftarrow_s \text{Enc}(\text{pk}^*, m^*)$	$c^* \leftarrow_s \text{Enc}(\text{pk}^*, m^*)$
$(\text{pk}, R, c) \leftarrow_s \mathcal{A}_2^{\text{O}_2}(\mathcal{M}, \text{pk}^*, s, c^*)$	$(\text{pk}, R, c) \leftarrow_s \mathcal{A}_2^{\text{O}_2}(\mathcal{M}, \text{pk}^*, s, c^*)$
return 1 iff $\exists(m, r)$ s.t.	return 1 iff $\exists(m, r)$ s.t.
$(c = \text{Enc}(\text{pk}, m; r)) \wedge$	$(c = \text{Enc}(\text{pk}, m; r)) \wedge$
$(c \neq c^* \vee \text{pk} \neq \text{pk}^*) \wedge$	$(c \neq c^* \vee \text{pk} \neq \text{pk}^*) \wedge$
$(m \neq \perp) \wedge R(m, m^*, \text{pk}, \text{pk}^*, c)$	$(m \neq \perp) \wedge R(m, \tilde{m}, \text{pk}, \text{pk}^*, c)$

In the experiments above

$$\begin{aligned} \text{if } \text{atk} = \text{cpa} \text{ then } \mathcal{O}_1 &= \epsilon & \text{and } \mathcal{O}_2 &= \epsilon, \\ \text{if } \text{atk} = \text{cca1} \text{ then } \mathcal{O}_1 &= \text{Dec}(\text{sk}^*, \cdot) & \text{and } \mathcal{O}_2 &= \epsilon, \\ \text{if } \text{atk} = \text{cca2} \text{ then } \mathcal{O}_1 &= \text{Dec}(\text{sk}^*, \cdot) & \text{and } \mathcal{O}_2 &= \text{Dec}^{(c^*)}(\text{sk}^*, \cdot), \end{aligned}$$

where $\text{Dec}^{(c)}(\text{sk}, \cdot)$ means that \mathcal{A} is allowed to query Dec oracle for any ciphertext distinct from the challenge ciphertext c^* .

2.3 Secret-Key Encryption

A *secret-key encryption* scheme (SKE) consists of a triple of algorithm $(\text{Gen}, \text{Enc}, \text{Dec})$ together with a message space \mathcal{M} and a key space \mathcal{K} , in which:

- The key generation algorithm Gen takes as input 1^λ and outputs a secret key $K \in \mathcal{K}$.

- The encryption algorithm Enc takes as input the secret key $K \in \mathcal{K}$ and a message $m \in \mathcal{M}$, and outputs a ciphertext c .
- The decryption algorithm Dec takes as input the secret key $K \in \mathcal{K}$ and a ciphertext c , and outputs a message $m \in \mathcal{M}$, or \perp denoting failure.

Let us consider the flavour of correctness needed for our scopes.

Definition 5 (ϵ -correctness). *A SKE scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is ϵ -correct if*

$$\Pr [\text{Dec}(\text{sk}, c) \neq m \mid K \leftarrow_s \text{Gen}(1^\lambda); c \leftarrow_s \text{Enc}(K, m)] \leq \epsilon.$$

To make our security proofs go through, we use the non-malleability security definition NM-ATK with $\text{ATK} \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$ introduced by Bellare et al. [2]. As highlighted in [4, 5] the original definitions were introduced in the asymmetric setting [6, 9, 19] but can be “lifted” to the symmetric setting using the encryption oracle based template of [2]. Hence, by leveraging the results of Bellare et al [6], NM-ATK for $\text{ATK} \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$ is equivalent to the indistinguishability-based counterpart IND-ATK [6, 4, 5].

Definition 6 (NM-CPA, NM-CCA1, NM-CCA2). *Given a set of relations \mathcal{R} , a SKE scheme Π is NM-ATK secure with respect to any relation $R \in \mathcal{R}$, if for any NM-ATK adversary $A = (A_0, A_1)$,*

$$\text{SKE}_{\Pi, A}^{\text{nm-atk}}(\lambda) \approx_c \text{SKE}_{\Pi, A, \$}^{\text{nm-atk}}(\lambda),$$

where the experiments are defined as follows

<i>Experiment</i> $\text{SKE}_{\Pi, A}^{\text{nm-atk}}(\lambda)$	<i>Experiment</i> $\text{SKE}_{\Pi, A, \$}^{\text{nm-atk}}(\lambda)$
$K^* \leftarrow_s \mathcal{K}$	$K^* \leftarrow_s \mathcal{K}$
$(\mathcal{M}, s) \leftarrow_s A_0^{O_1}(1^\lambda)$	$(\mathcal{M}, s) \leftarrow_s A_0^{O_1}(1^\lambda)$
$m^* \leftarrow_s \mathcal{M}$	$m^*, \tilde{m} \leftarrow_s \mathcal{M}$
$c^* \leftarrow_s \text{Enc}(K^*, m^*)$	$c^* \leftarrow_s \text{Enc}(K^*, m^*)$
$(R, c') \leftarrow_s A_1^{O_2}(\mathcal{M}, s, c^*)$	$(R, c') \leftarrow_s A_1^{O_2}(\mathcal{M}, s, c^*)$
$m' := \text{Dec}(K^*, c')$	$m' := \text{Dec}(K^*, c')$
return 1 iff	return 1 iff
$(m' \neq \perp) \wedge (c' \neq c^*) \wedge R(m^*, m')$	$(m' \neq \perp) \wedge (c' \neq c^*) \wedge R(\tilde{m}, m')$

In the experiments above

$$\begin{aligned}
&\text{if } \text{atk} = \text{cpa} \text{ then } O_1 = \text{Enc}(K^*, \cdot) \quad \text{and } O_2 = \text{Enc}(K^*, \cdot), \\
&\text{if } \text{atk} = \text{cca1} \text{ then } O_1 = \text{Enc}(K^*, \cdot), \text{Dec}(K^*, \cdot) \text{ and } O_2 = \text{Enc}(K^*, \cdot), \\
&\text{if } \text{atk} = \text{cca2} \text{ then } O_1 = \text{Enc}(K^*, \cdot), \text{Dec}(K^*, \cdot) \text{ and } O_2 = \text{Enc}(K^*, \cdot), \text{Dec}^{(c^*)}(K^*, \cdot),
\end{aligned}$$

where $\text{Dec}^{(c^*)}(K^*, \cdot)$ means that A is allowed to query the Dec oracle for any ciphertext c distinct from the challenge ciphertext c^* .

3 Completely Non-Malleable KEMs

In this section, we formally introduce key encapsulation methods (KEMs). Then, following the work on completely non-malleable PKE schemes of Fischlin [12], and then Ventre and Visconti [22], we extend the notions of complete non-malleability of PKEs to the context of KEMs. We follow a similar blueprint of [22] (that, in turn, bases its definitions on [3]) by introducing three indistinguishability-based security notions for completely non-malleable KEMs dubbed NM-CPA*, NM-CCA1* and NM-CCA2*. We further introduce,

in [Section 3.1](#), three simulation-based notions dubbed SNM-CPA*, SNM-CCA1* and SNM-CCA2* and investigate the relationship between indistinguishability-based and the simulation-based notions.

A KEM scheme consists of a triple of algorithms $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$, together with a key space \mathcal{K} , in which:

- The key generation algorithm Gen takes as input 1^λ and outputs a public-private key pair (pk, sk) .
- The encapsulation algorithm Encaps takes as input a public key pk , and output a ciphertext c as well as a key K .
- The decapsulation algorithm Decaps takes as input a private key sk and a ciphertext c and returns a key $K \in \mathcal{K}$ or \perp (denoting failure).

First of all, we start by considering the flavour of correctness needed for our scopes.

Definition 7 (ϵ -correctness). *A KEM scheme $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$ is ϵ -correct if*

$$\Pr [\text{Decaps}(\text{sk}, c) \neq K \mid (\text{pk}, \text{sk}) \leftarrow_{\$} \text{Gen}(1^\lambda); (c, K) \leftarrow_{\$} \text{Encaps}(\text{pk})] \leq \epsilon.$$

Before diving into the definitions complete non-malleability for KEMs, we introduce the notion of complete relation for KEMs, firstly defined by Fischlin [12] in the setting of completely non-malleable PKE. A complete relation R in the KEM setting is a probabilistic algorithm taking as inputs two public keys pk and pk' , two encapsulation keys K and K^* , and a ciphertext c' . It outputs 1 if the relation is satisfied, and 0 otherwise. We will refer as \mathcal{R} to be the set of complete relations.

In the indistinguishability-based notion of complete non-malleability, we ask the adversary to distinguish between two experiments. In both of them, we let the adversary learn the challenge public key, ciphertext and encapsulated key, and then make the adversary output a new public key pk' , a relation R and a new ciphertext c' . If there exists a key $K' \neq K$ and randomness r such that c' and K' can be obtained by running the encapsulation algorithm with $\text{pk}' \neq \text{pk}^*$ and randomness r , then the experiment will output 1. In the left-side experiment the adversary will receive the key K^* encapsulated in c^* , while in the right-side experiment the key K^* is sampled randomly and hence totally unrelated from the key encapsulated in the received ciphertext. Note that the adversary may come up with a triple (pk', R, c') such that there exists a key encapsulated in c' satisfying the relation $R(K', K^*, \text{pk}', \text{pk}^*, c')$, but the adversary may have negligible advantage in distinguishing whether the key was encapsulated in c^* or it was randomly chosen.

Definition 8 (NM-CPA*, NM-CCA1*, NM-CCA2*). *Given a set of relations \mathcal{R} , a key-encapsulation mechanism $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$ is NM-ATK* secure with respect to any relation $R \in \mathcal{R}$, if for any NM-ATK* adversary $A = (A_1, A_2)$ for $\text{atk} \in \{\text{cpa}, \text{cca1}, \text{cca2}\}$ and for all $\lambda \in \mathbb{N}$,*

$$\text{KEM}_{\Pi, A}^{\text{nm-atk}^*}(\lambda) \approx_c \text{KEM}_{\Pi, A, \$}^{\text{nm-atk}^*}(\lambda),$$

where the experiments are defined as follows:

<i>Experiment</i> $\text{KEM}_{\Pi, A}^{\text{nm-atk}^*}(\lambda)$	<i>Experiment</i> $\text{KEM}_{\Pi, A, \$}^{\text{nm-atk}^*}(\lambda)$
$(\text{pk}^*, \text{sk}^*) \leftarrow_{\$} \text{Gen}(1^\lambda)$	$(\text{pk}, \text{sk}) \leftarrow_{\$} \text{Gen}(1^\lambda)$
$\text{st} \leftarrow_{\$} A_1^{\text{O}1}(\text{pk}^*)$	$\text{st} \leftarrow_{\$} A_1^{\text{O}1}(\text{pk}^*)$
$(c^*, K^*) \leftarrow_{\$} \text{Encaps}(\text{pk}^*)$	$K^* \leftarrow_{\$} \{0, 1\}^\lambda$
$(\text{pk}', R, c') \leftarrow_{\$} A_2^{\text{O}2}(\text{pk}^*, c^*, K^*, \text{st})$	$(\hat{c}, \hat{K}) \leftarrow_{\$} \text{Encaps}(\text{pk}^*)$
return 1 if $\exists(K', r)$ such that	$(\text{pk}', R, c') \leftarrow_{\$} A_2^{\text{O}2}(\text{pk}^*, \hat{c}, K^*, \text{st})$
$((c', K') = \text{Encaps}(\text{pk}'; r)) \wedge$	return 1 if $\exists(K', r)$ such that
$(\text{pk}' \neq \text{pk}^* \vee K' \neq K^*) \wedge (K' \neq \perp)$	$((c', K') = \text{Encaps}(\text{pk}'; r)) \wedge$
$\wedge R(K', K^*, \text{pk}', \text{pk}^*, c')$	$(\text{pk}' \neq \text{pk}^* \vee K' \neq K^*) \wedge (K' \neq \perp)$
	$\wedge R(K', K^*, \text{pk}', \text{pk}^*, c')$

In the experiments above,

$$\begin{aligned}
& \text{if } \text{atk} = \text{cpa} \text{ then } O_1 = \epsilon \quad \text{and } O_2 = \epsilon, \\
& \text{if } \text{atk} = \text{cca1} \text{ then } O_1 = \text{Decaps}(\text{sk}^*, \cdot) \text{ and } O_2 = \epsilon, \\
& \text{if } \text{atk} = \text{cca2} \text{ then } O_1 = \text{Decaps}(\text{sk}^*, \cdot) \text{ and } O_2 = \text{Decaps}^{(c^*)}(\text{sk}^*, \cdot),
\end{aligned}$$

where $\text{Decaps}^{(c^*)}(\text{sk}, \cdot)$ means that A is allowed to query Decaps algorithm for any ciphertext distinct from the challenge ciphertext c^* .

In [12], Fischlin showed that there exist encryption and signatures schemes that are not NM-CCA2* secure, even though they are NM-CCA2 secure, i.e. the CCA secure in standard non-malleability notion of PKE (see [6] for the formal definition). As we show in the following theorem, this holds also in the case of KEMs. Let $\text{ATK} \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$.

Theorem 1. *Assume that there exists a NM-ATK secure KEM $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$; then there exists a NM-ATK secure KEM $\Pi' = (\text{Gen}', \text{Encaps}', \text{Decaps}')$ which is not NM-ATK* secure.*

Proof. The intuition behind the proof is to define the scheme Π' in such a way that the adversary can leverage its structure to succeed in the NM-ATK* experiment. In particular, we define Π' in the following way:

Algorithm $\text{Gen}'(1^\lambda)$	Algorithm $\text{Encaps}'(\text{pk}')$	Algorithm $\text{Decaps}'(\text{sk}, c)$
$(\text{pk}, \text{sk}) = \text{Gen}(1^\lambda)$	Parse pk' as $\text{pk}' := \text{pk} b$	$K := \text{Decaps}'(\text{sk}, c)$
$b \leftarrow_s \{0, 1\}$	$(c, K) \leftarrow_s \text{Encaps}(\text{pk})$	
$\text{pk}' := \text{pk} b$		
return (pk', sk)	return (c, K)	return K

We can clearly see that Π' is not NM-ATK* secure. Indeed, an efficient adversary A receiving a public key pk , the challenge ciphertext c and an encapsulation key K (either the real or the fake one), just have to flip the last bit of pk denoted $\text{pk}' := \text{pk}||\bar{b}$, a relation R and the challenge ciphertext c . In this case, A will always succeeds in breaking NM-ATK* security of Π' .

However, Π' is still NM-ATK secure. Indeed, this is true because the adversary has to break NM-ATK security of Π' under the key pk , but for how Π' is defined, this is equivalent to break NM-ATK security of Π . If an adversary is able to break NM-ATK security of Π' , then he can also break NM-ATK security of Π , and this represent a contradiction to our assumption that Π is NM-ATK secure. \square

3.1 Relation between NM-ATK* and SNM-ATK*

We introduce the simulation-based variant of NM-ATK* security, dubbed SNM-ATK*. The left-side of the simulation-based experiment is identical to the left-side of the indistinguishability-based notion (Definition 8), whereas in the right-side experiment the challenge key K^* is sampled at random, and the adversary is replaced with a simulator knowing only public information, i.e. the challenge public key pk^* .

Definition 9 (SNM-CPA*, SNM-CCA1*, SNM-CCA2*). *Given a set of relations \mathcal{R} , a key-encapsulation mechanism $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$ is secure in the sense of SNM-ATK* with respect to any relation $R \in \mathcal{R}$, if for any SNM-ATK* adversary $A = (A_1, A_2)$, for $\text{atk} \in \{\text{cpa}, \text{cca1}, \text{cca2}\}$, there exists a simulator S such that*

$$\text{KEM}_{\Pi, A}^{\text{snm-atk}^*}(\lambda) \approx_c \text{KEM}_{\Pi, S}^{\text{snm-atk}^*}(\lambda),$$

where the experiments $\text{KEM}_{\Pi, A}^{\text{snm-atk}^*}(\lambda)$ and $\text{KEM}_{\Pi, S}^{\text{snm-atk}^*}(\lambda)$ are defined as follows:

<p>Experiment $\text{KEM}_{\Pi, A}^{\text{snm-atk}^*}(\lambda)$</p> <hr style="border: 0.5px solid black;"/> <p>$(\text{pk}^*, \text{sk}^*) \leftarrow_{\\$} \text{Gen}(1^\lambda)$ $\text{st} \leftarrow_{\\$} A_1^{\text{O}_1}(\text{pk}^*)$ $(c^*, K^*) \leftarrow_{\\$} \text{Encaps}(\text{pk}^*)$</p> <p>$(\text{pk}', R, c') \leftarrow_{\\$} A_2^{\text{O}_2}(\text{pk}^*, c^*, K^*, \text{st})$ return 1 if $\exists(K', r)$ such that $((c', K') = \text{Encaps}(\text{pk}'; r)) \wedge$ $(\text{pk}' \neq \text{pk} \vee K' \neq K^*) \wedge (K' \neq \perp)$ $\wedge R(K', K^*, \text{pk}', \text{pk}^*, c')$</p>	<p>Experiment $\text{KEM}_{\Pi, S}^{\text{snm-atk}^*}(\lambda)$</p> <hr style="border: 0.5px solid black;"/> <p>$(\text{pk}^*, \text{sk}^*) \leftarrow_{\\$} \text{Gen}(1^\lambda)$ $K^* \leftarrow_{\\$} \{0, 1\}^\lambda$</p> <p>$(\text{pk}', R, c') \leftarrow_{\\$} S(\text{pk}^*)$ return 1 if $\exists(K', r)$ such that $((c', K') = \text{Encaps}(\text{pk}'; r)) \wedge$ $(\text{pk}' \neq \text{pk} \vee K' \neq K^*) \wedge (K' \neq \perp)$ $\wedge R(K', K^*, \text{pk}', \text{pk}^*, c')$</p>
--	--

In the experiments above,

if $\text{atk} = \text{cpa}$ then $\text{O}_1 = \epsilon$ and $\text{O}_2 = \epsilon$,
 if $\text{atk} = \text{cca1}$ then $\text{O}_1 = \text{Decaps}(\text{sk}, \cdot)$ and $\text{O}_2 = \epsilon$,
 if $\text{atk} = \text{cca2}$ then $\text{O}_1 = \text{Decaps}(\text{sk}, \cdot)$ and $\text{O}_2 = \text{Decaps}^{(c^*)}(\text{sk}, \cdot)$,

where $\text{Decaps}^{(c^*)}(\text{sk}, \cdot)$ means that A is allowed to query Decaps algorithm for any ciphertext distinct from the challenge ciphertext c^* .

We will now show the relationship between the indistinguishability-based definitions and simulation-based ones.

In particular, we will show that our indistinguishability-based definition of NM-CPA* security implies the corresponding simulation-based definition, and then argue about the restrictions we need to prove equivalence between NM-CCA1*/NM-CCA2* and the simulation-based counterpart. Conversely, we show that the simulation-based definition NM-ATK* implies the indistinguishability-based definition SNM-ATK* for $\text{ATK} \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$.

Theorem 2 (NM-CPA* \implies SNM-CPA*). *If a KEM $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$ is NM-CPA* secure (Definition 8), then Π is SNM-CPA* secure (Definition 9).*

Proof. We follow a similar approach of Ventre and Visconti [22] to prove that a if PKE is NM-CPA*, then it is also secure for SNM-CPA*. The intuition behind the proof is that, given an adversary $A = (A_1, A_2)$, the simulator S simply runs A . In particular, we define S as follows:

Simulator $S(\text{pk})$:

$\text{st} \leftarrow_{\$} A_1(\text{pk})$
 $(\hat{c}, \hat{K}) \leftarrow_{\$} \text{Encaps}(\text{pk})$
 $(\text{pk}', R, c') \leftarrow_{\$} A_2(\text{pk}, \hat{c}, \hat{K}, \text{st})$
return (pk', R, c')

The simulator can directly run the adversary A because A doesn't have access to any oracle, which in turn means that S doesn't need to know the private key corresponding to the challenge public key pk . In order to show that Π is secure in the sense of SNM-CPA*, we first note that the experiment $\text{KEM}_{\Pi, A}^{\text{nm-cpa}^*}(\lambda)$ is indeed equivalent to $\text{KEM}_{\Pi, A}^{\text{snm-cpa}^*}(\lambda)$, and thus,

$$\text{KEM}_{\Pi, A}^{\text{nm-cpa}^*}(\lambda) \equiv \text{KEM}_{\Pi, A}^{\text{snm-cpa}^*}(\lambda).$$

Now, by replacing the simulator in $\text{KEM}_{\Pi, S}^{\text{snm-cpa}^*}$ with the simulator given above, we get that the resulting experiment is equivalent to that of $\text{KEM}_{\Pi, A, S}^{\text{nm-cpa}^*}$. Thus, we get that

$$\text{KEM}_{\Pi, \mathcal{A}, \mathcal{S}}^{\text{nm-cpa}^*}(\lambda) \equiv \text{KEM}_{\Pi, \mathcal{S}}^{\text{snm-cpa}^*}(\lambda).$$

□

Extending to NM-CCA1* and NM-CCA2* security. Unfortunately, the direct approach used to prove [Theorem 3](#) cannot be used to show that the game-based definitions of NM-CCA1* and NM-CCA2* imply the corresponding simulation-based counterpart. A NM-ATK* adversary \mathcal{A} , for $\text{ATK} \in \{\text{CCA1}, \text{CCA2}\}$, has access to the decapsulation oracle, and therefore the simulator \mathcal{S} in order to perfectly simulate the NM-ATK* experiment, must be able to simulate the queries to such oracle. However, \mathcal{S} doesn't have access to any oracle, so the only way for the simulator to simulate such queries is to know the secret key corresponding to the given public key pk . Naturally, this cannot be a realistic assumption, and therefore we let \mathcal{S} to generate its own pair of public-private keys $(\hat{\text{pk}}, \hat{\text{sk}})$, and use $\hat{\text{sk}}$ to answer the queries to the decapsulation oracle made by \mathcal{A} . This assumption has an effect on the set of relations that can consider. Since the simulator sends its own public key $\hat{\text{pk}}$ to the adversary ignoring the challenge key pk^* , the relations under consideration must be independent of pk^* as well, otherwise \mathcal{A} has no way to succeed in the SNM-ATK* experiment without knowing the challenge public key pk^* . As introduced by Ventre and Visconti in [\[22\]](#), we will consider such a special type of relations, called *lacking* relations. More formally, a relation R is lacking if it is a complete relation that ignores the input of the challenge public key pk^* , i.e. R is lacking if and only if, for all $K^*, K' \in \mathcal{K}$, all ciphertext c' produced by the Encaps algorithm, all public keys pk^*, pk' produced by the key generation algorithm, $R(K', K^*, \text{pk}', \cdot, c') = R(K', K^*, \text{pk}', \text{pk}^*, c')$.

Theorem 3 (NM-ATK* \implies SNM-ATK*). *Let \mathcal{R} denote the set of lacking relations, and $\text{ATK} \in \{\text{CCA1}, \text{CCA2}\}$. If a KEM scheme $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$ is NM-ATK* secure ([Definition 8](#)) with respect to \mathcal{R} then Π is SNM-ATK* secure with respect to \mathcal{R} ([Definition 9](#)).*

Due to its similarity to the proof of [Theorem 2](#), we have provided the sketch of the proof.

Proof. [Sketch] When the simulator \mathcal{S} receives the challenge public key pk , generates a new pair of keys and runs an adversary \mathcal{A} in the simulation-based experiment, on input the new public key. \mathcal{S} computes the encapsulation using the new public key and gives the resulting ciphertext and key to \mathcal{A} as the challenge ciphertext and key. Next, the simulator uses the corresponding encapsulation key to answer the queries to the decapsulation oracle made by \mathcal{A} . At the end, when the adversary returns a public key pk' , a relation R , and a ciphertext c' , the simulator just returns (pk', R, c') . Since \mathcal{A} is a NM-ATK* adversary, the final ciphertext c' is obtained using a public key that is different from the given one. Furthermore, since we are considering lacking relations, it is straightforward to see that \mathcal{S} succeeds whenever \mathcal{A} does. □

We will now show that the converse implication holds for $\text{ATK} \in \{\text{CPA}, \text{CCA1}, \text{CCA2}\}$.

Theorem 4 (SNM-ATK* \implies NM-ATK*). *If a KEM scheme $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$ is SNM-ATK* secure ([Definition 9](#)) with respect to a set of relations \mathcal{R} then Π is NM-ATK* secure with respect to \mathcal{R} ([Definition 8](#)).*

Proof. The converse implication follows directly. Let us consider an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ of the NM-ATK* experiment as follows.

$$\begin{array}{l|l} \hline \text{Adversary } \mathcal{A}_1^{\text{O}_1}(\text{pk}^*): & \text{Adversary } \mathcal{A}_2^{\text{O}_2}(\text{pk}^*, c^*, K^*, \text{st}): \\ \text{return st} := \epsilon & (\text{pk}', R, c') \leftarrow \mathcal{S}(\text{pk}^*) \\ & \text{return } (\text{pk}', R, c') \\ \hline \end{array}$$

In order to show that Π is secure in the sense of NM-CPA*, we first note that the experiment $\text{KEM}_{\Pi, \mathcal{A}}^{\text{snm-atk}^*}(\lambda)$ is indeed equivalent to $\text{KEM}_{\Pi, \mathcal{A}}^{\text{nm-atk}^*}(\lambda)$, and thus,

$$\text{KEM}_{\Pi, \mathcal{A}}^{\text{nm-atk}^*}(\lambda) \equiv \text{KEM}_{\Pi, \mathcal{A}}^{\text{snm-atk}^*}(\lambda).$$

Now, by replacing the adversary A in $\text{KEM}_{\Pi, A, \$}^{nm-atk^*}(\lambda)$ with the one described above, it is straightforward to see that

$$\text{KEM}_{\Pi, A, \$}^{nm-atk^*}(\lambda) \equiv \text{KEM}_{\Pi, S}^{snm-atk^*}(\lambda).$$

□

4 Analysis of Fujisaki-Okamoto transforms

In the following, we analyze complete non-malleability of the FO transforms. To do that, we will consider the modular treatment of the FO transforms pursued by [14]. Each FO transform is an application of two transformations, namely T and U . T takes as input an IND-CPA/OW-CPA secure PKE scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ and a random oracle G , and outputs a deterministic OW-PCVA PKE scheme $\Pi_1 = (\text{Gen}_1, \text{Enc}_1, \text{Dec}_1)$ (when Π is OW-CPA, it should also satisfy γ -uniformity for a sufficiently large γ). The encryption algorithm of the transformed PKE runs the encryption algorithm of the underlying IND-CPA secure PKE scheme Π but uses $G(m)$ as its randomness, i.e. $\text{Enc}_1(\text{pk}, m) := \text{Enc}(\text{pk}, m; G(m))$. The decryption algorithm runs the decryption algorithm of the underlying PKE scheme, i.e. $m' := \text{Dec}(\text{sk}, c)$, and returns \perp if $m' = \perp$ or the re-encryption of m under public-key pk and randomness $G(m)$ does not match with c , i.e. if $\text{Enc}(\text{pk}, m'; G(m')) \neq c$.

Given a PKE scheme satisfying OW-PCVA or some different flavour of one-wayness (depending on the transformation we are going to use) and a random oracle H , four variants of the transformation U can be used to produce an IND-CCA2 KEM scheme. In Fig. 1 we recall the algorithms needed to instantiate the U^\perp and U^\times transformations. The algorithms for the U_m^\perp and U_m^\times transformations are the same as U^\perp and U^\times respectively, except that the encapsulation algorithm computes the key K as $H(m)$, and the decapsulation algorithm outputs $K := H(m)$ when $m \neq \perp$.

<p>Algorithm $\text{Gen}^\times(1^\lambda)$</p> <hr/> <p>$(\text{pk}', \text{sk}') \leftarrow_{\\$} \text{Gen}_1(1^\lambda)$ $s \leftarrow_{\\$} \mathcal{M}$ $\text{sk} := (\text{sk}', s)$ return (pk', sk)</p> <p>Algorithm $\text{Decaps}^\perp(\text{sk}, c)$</p> <hr/> <p>$m' := \text{Dec}_1(\text{sk}, c)$ if $m' = \perp$ return \perp return $K := H(m', c)$</p>	<p>Algorithm $\text{Encaps}(\text{pk})$</p> <hr/> <p>$m \leftarrow_{\\$} \mathcal{M}$ $c \leftarrow_{\\$} \text{Enc}_1(\text{pk}, m)$ $K := H(m, c)$ return (K, c)</p> <p>Algorithm $\text{Decaps}^\times(\text{sk}, c)$</p> <hr/> <p>Parse $\text{sk} = (\text{sk}', s)$ $m' := \text{Dec}_1(\text{sk}', c)$ if $m' = \perp$ return $K := H(m', s)$ return $K := H(m', c)$</p>
--	---

Figure 1: Algorithms needed by the transformations $U^\perp[\Pi_1, H] = (\text{Gen}_1, \text{Encaps}, \text{Decaps}^\perp)$ and $U^\times[\Pi_1, H] = (\text{Gen}^\perp, \text{Encaps}, \text{Decaps}^\times)$.

In Section 4.1, we will analyze the U_m^\perp/U_m^\times transformations and show that they are not completely non-malleable by giving an attack that can be performed when the underlying OW-PCVA PKE scheme is obtained by applying the transformation T to a widely known IND-CPA PKE scheme.

Then, in Section 4.2, we will analyze the U^\perp/U^\times transformations and show that they are not completely non-malleable by constructing a contrived OW-PCVA PKE scheme that, when given as input to U^\perp/U^\times , leads to a KEM whose NM-ATK* security can be easily broken. Then, we show that it suffices to assume a very natural property of the underlying OW-PCVA PKE scheme, named public-key uniqueness, to achieve complete non-malleability.

Finally, in [Section 4.3](#), we show that it is possible to achieve a completely non-malleable KEM without assuming public-key uniqueness of the underlying PKE scheme with a little tweak to U^\perp/U^\times .

4.1 Analysis of the $U_m^{\perp/\times}$ transformations

In the following, we will show that U_m^\perp and U_m^\times transformations lead to a KEM that is not completely non-malleable. In particular, we show a concrete attack that can be carried out to both $\tilde{\Pi}_m^\perp := U_m^\perp[\Pi, \mathbb{G}, \mathbb{H}]$ and $\tilde{\Pi}_m^\times := U_m^\times[\Pi, \mathbb{G}, \mathbb{H}]$ even against the weaker notion of NM-CPA*. Let Π be the El-Gamal encryption scheme.

When running the experiment $\text{KEM}_{\tilde{\Pi}_m^\perp, \mathbb{A}}^{nm\text{-cpa}^*}(\lambda)$ with $\tilde{\Pi}_m^\perp$ (resp. $\tilde{\Pi}_m^\times$), an efficient adversary \mathbb{A} receives as input a public key $\text{pk}^* = (\text{params}, h)$, the challenge ciphertext c^* , and an encapsulation key \mathbb{K}^* . The challenge ciphertext c^* is computed as $(c_1^* = g^r, c_2^* = h^r \cdot m^*)$ with $r = \mathbb{G}(m^*)$, $\text{params} = (\mathbb{G}, g, q)$, where g is the generator of a cyclic group \mathbb{G} of order q , $h = g^x$ for $x \leftarrow_s \mathbb{Z}_q$, and m^* is a randomly sampled message from \mathbb{G} . The challenge key \mathbb{K}^* is either the output of the encapsulation algorithm on input m^* (i.e., $\mathbb{K}^* = \mathbb{H}(m^*)$), or it is sampled from the uniform distribution over the key space \mathcal{K} .

Now, \mathbb{A} can craft a new public key $\text{pk}' = (\text{params}, h \cdot g^{x'})$, for $x' \leftarrow_s \mathbb{Z}_q$, and a ciphertext $c' = (c_1^*, c_2^* \cdot c_1^{*x'})$. It is straightforward to see that, since $\text{Enc}(\text{pk}', m') = (g^r, g^{(x+x')r} m^*) = c'$, there exist a key \mathbb{K}' such that $(c', \mathbb{K}') = \text{Encaps}(\text{pk}')$, where $c' = \text{Enc}(\text{pk}', m^*)$. Finally, since $\mathbb{K}' := \mathbb{H}(m')$ with $m' = m^*$, then $\mathbb{K}' = \mathbb{K}^*$.

Now, \mathbb{A} can define the relation between pk^* and pk' as follows:

$$R_{\text{pk}}(\text{pk}^* = (\text{params}, h), \text{pk}' = (\text{params}, h')) = 1 \text{ iff } h' = h \cdot g^{x'}, x' \in \mathbb{Z}_p.$$

Moreover, \mathbb{A} can define the relation between \mathbb{K}^* and \mathbb{K}' as the identity relation, i.e. $R_{\mathbb{K}}(\mathbb{K}^*, \mathbb{K}') = 1$ iff $\mathbb{K}^* = \mathbb{K}'$. As shown above, \mathbb{A} can come up with a $\text{pk}' \neq \text{pk}$ satisfying the relation $R(\mathbb{K}^*, \mathbb{K}', \text{pk}^*, \text{pk}, c') = R_{\mathbb{K}}(\mathbb{K}^*, \mathbb{K}') \wedge R_{\text{pk}}(\text{pk}^*, \text{pk}')$. Such relation, when the ciphertext c' is computed as above (i.e. $c' = (c_1^*, c_2^* \cdot c_1^{*x'})$), leads the $\text{KEM}_{\tilde{\Pi}_m^\perp, \mathbb{A}}^{nm\text{-cpa}^*}(\lambda)$ experiment output 1 and the $\text{KEM}_{\tilde{\Pi}_m^\times, \mathbb{A}, \mathbb{S}}^{nm\text{-cpa}^*}(\lambda)$ output 0 with non-negligible probability.

In the following, we will show that the U^\perp/U^\times transformations can be turned into a completely non-malleable KEM without much effort.

4.2 Analysis of the $U^{\perp/\times}$ transformations

The U^\perp and U^\times transformations do not naturally satisfy complete non-malleability. Indeed, given an OW-PCVA PKE scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$, we can construct a contrived OW-PCVA PKE scheme $\Pi' = (\text{Gen}', \text{Enc}', \text{Dec}')$ that is trivially insecure against the NM-CPA* experiment. The scheme Π' follows:

Algorithm $\text{Gen}'(1^\lambda)$	Algorithm $\text{Enc}'(\text{pk}', m)$	Algorithm $\text{Dec}'(\text{sk}, m)$
$(\text{pk}, \text{sk}) \leftarrow_s \text{Gen}(1^\lambda)$	Parse $\text{pk}' = (\text{pk}, b)$	$m := \text{Dec}(\text{sk}, m)$
$b \leftarrow_s \{0, 1\}$	$c \leftarrow_s \text{Enc}(\text{pk}, m)$	return m
$\text{pk}' := \text{pk} \ b$	return c	
return (pk', sk)		

It is straightforward to see that the adversary of NM-CPA*, when run with $\tilde{\Pi}^\perp := U^\perp[\Pi', \mathbb{H}]$, distinguishes between the two experiments with non-negligible probability. Indeed, an adversary \mathbb{A} can come up with a public key pk' different of the public key generated by $\tilde{\Pi}^\perp$ that leads to the same ciphertext c^* encapsulating \mathbb{K}^* as computed in the $\text{KEM}_{\tilde{\Pi}^\perp, \mathbb{A}}^{nm\text{-cpa}^*}$ experiment. To be precise, the adversary \mathbb{A} might define a relation $R(\mathbb{K}^*, \mathbb{K}', \text{pk}^*, \text{pk}, \cdot) = R_{\mathbb{K}}(\mathbb{K}^*, \mathbb{K}') \wedge R_{\text{pk}}(\text{pk}^*, \text{pk}')$, where $R_{\mathbb{K}}(\mathbb{K}^*, \mathbb{K}')$ is the identity function and $R_{\text{pk}}(\text{pk}^*, \text{pk}') = 1$ iff, given that pk^* can be parsed as (pk, b) with $b \in \{0, 1\}$, then pk' equals $(\text{pk}, 1 - b)$. If c^* is the encryption of a message m^* under pk^* , then it trivially encrypts m^* under pk' . Hence, the key $\mathbb{K}^* := \mathbb{H}(m^*, c^*)$ will be the same.

An idea to avoid such an artificial attack, is to restrict the set of admitted PKE schemes to the ones for which such attack is not possible to carry out in the first place. As we will show, this suffice to guarantee complete non-malleability of $\tilde{\Pi}^\perp$. To do that, we introduce a natural property called public-key uniqueness, informally stating that it is infeasible for an adversary to come up with two different public keys leading to the same ciphertext when encrypting the same message. For example, in El-Gamal, the encryption of the same message with two different public keys will always lead to a different ciphertext for any possible random coins and any possible message. This property, which we formalize below, is indeed achieved by most of the known PKE schemes.

Definition 10 (Public-Key Uniqueness). *A public-key encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ with message space \mathcal{M} is public-key unique if, for all unbounded/PPT adversary A , there exists a negligible function $\text{negl}(\lambda)$ such that*

$$\Pr \left[\exists r, r', m : \begin{array}{l} \text{Enc}(\text{pk}, m; r) = \text{Enc}(\text{pk}', m; r') \\ \wedge (\text{pk} \neq \text{pk}') \end{array} \mid (\text{pk}, \text{pk}') \leftarrow_s A(1^\lambda) \right] \leq \text{negl}(\lambda).$$

We say that Π is perfect public-key unique if do not exist two public keys pk and pk' with $\text{pk}' \neq \text{pk}$ such that $\text{Enc}(\text{pk}, m; r) = \text{Enc}(\text{pk}', m; r')$ for any $m \in \mathcal{M}$ and any randomness r, r' .

In the following, we show that the KEM $\tilde{\Pi}^\perp = \text{U}^\perp[\Pi_1, \text{H}]$ is completely non-malleable when the PKE Π_1 is OW-PCVA and satisfies perfect public-key uniqueness. Note that our result trivially extends when Π_1 satisfies public-key uniqueness against PPT or unbounded adversaries.

Theorem 5 (Π_1 pk-unique OW-PCVA $\xrightarrow{\text{ROM}}$ $\tilde{\Pi}^\perp$ NM-CCA2*). *Assuming the existence of a random oracle H , if Π_1 is a correct OW-PCVA secure PKE (Definition 3) satisfying perfect public-key uniqueness (Definition 10), then $\tilde{\Pi}^\perp$ defined as above is a correct NM-CCA2* secure KEM (Definition 8).*

Proof. Correctness of $\tilde{\Pi}^\perp$ trivially follows from the correctness of Π_1 . The idea behind the proof is to simulate the decapsulation oracle without using the secret key. We can do that by answering the decapsulation queries with a random key, and next simulate the random oracle H by using the plaintext checking oracle $\text{PCO}(\text{sk}^*, \cdot, \cdot)$, provided by the OW-PCVA game. Furthermore, we will use the ciphertext validity oracle $\text{CVO}^{(c^*)}(\text{sk}^*, \cdot)$ in order to reject decapsulation queries for invalid ciphertexts. Before proceeding with the proof we need to make the following observations. Since a NM-CCA2* adversary against $\tilde{\Pi}^\perp$ is allowed to choose both the public key and the ciphertext, the decapsulation oracle can receive ciphertexts encrypted using a public key $\text{pk}' \neq \text{pk}^*$. However, since the decapsulation oracle does not know the secret key associated to pk' , we cannot require to it to check the validity of such ciphertexts. In other words, the decapsulation oracle is only required to check the validity of ciphertexts encrypted using the challenge public key pk^* . A similar observation shall be done for the oracle H . Due to the fact that also H can receive ciphertexts encrypted using any public key distinct from pk^* , we should be able to check the validity of such ciphertexts. In other words, given a message m and a ciphertext c , we should allow H to check whether $m = \text{Dec}_1(\text{sk}', c)$. However, since H doesn't know the secret key associated with pk' , this desired behavior cannot be achieved. For this reason, the same approach used for $\text{Decaps}_\perp^{(c^*)}(\text{sk}^*, \cdot)$ applies. Let B be a PPT adversary that breaks NM-CCA2* security of $\tilde{\Pi}^\perp$ with non-negligible probability issuing a polynomially bounded number of queries to $\text{Decaps}_\perp^\perp$ and H . The sequence of games we are going to consider is described in Fig. 2.

Lemma 1. $\text{KEM}_{\tilde{\Pi}^\perp, B}^{nm-cca2^*}(\lambda) \equiv \mathbf{G}_0^B(\lambda)$.

Proof. Let us start by noticing that in game \mathbf{G}_0 the challenger takes a uniform message m^* , computes $c^* \leftarrow_s \text{Enc}_1(\text{pk}^*, m^*)$, $K^* := \text{H}(m^*, c^*)$ and outputs (pk^*, c^*, K^*) . This game coincides exactly with the left experiment of the NM-CCA2* definition. Thus, $\text{KEM}_{\tilde{\Pi}^\perp, B}^{nm-cca2^*}(\lambda) \equiv \mathbf{G}_0^B(\lambda)$. \square

Lemma 2. $\mathbf{G}_0^B(\lambda) \equiv \mathbf{G}_1^B(\lambda)$.

Games $G_0^B(\lambda) - G_3^B(\lambda)$

$(pk^*, sk^*) \leftarrow \text{Gen}_1(1^\lambda)$
 $st \leftarrow \text{B}_1^{\text{Decaps}_\perp(sk^*, \cdot), H(\cdot, \cdot)}(pk^*)$
 $m^* \leftarrow \mathcal{M}$
 $c^* \leftarrow \text{Enc}_1(pk, m^*)$
 $K^* := H(m^*, c^*) \quad \#G_0, G_1$
 $K^* \leftarrow \{0, 1\}^\lambda \quad \#G_2, G_3$
 $(pk', R, c') \leftarrow \text{B}_2^{\text{Decaps}_\perp^{(c^*)}(sk^*, \cdot), H(\cdot, \cdot)}(pk^*, c^*, K^*, st)$
return 1 iff $\exists(K', r)$ such that
 $(c', K') := \text{Encaps}(pk'; r) \wedge (K' \neq \perp)$
 $\wedge (pk' \neq pk^* \vee K' \neq K^*)$
 $\wedge R(K', K^*, pk', pk^*, c')$

Oracle $\text{Decaps}_\perp^{(c^*)}(sk^*, c) \quad \#G_0, G_3$

$m' := \text{Dec}_1(sk^*, c)$
if $m' = \perp$
 return \perp
return $K := H(m', c)$

Oracle $H(m, c)$

if $\exists K$ s.t. $(m, c, K) \in \mathcal{L}_H$
 return K
 $K \leftarrow \mathcal{K}$
if $\text{Dec}_1(sk^*, c) = m \quad \#G_1, G_2$
 if $c = c^* \quad \#G_2$
 return $\perp \quad \#G_2$
 if $\exists K'$ s.t. $(c, K') \in \mathcal{L}_D \quad \#G_1$
 $K := K' \quad \#G_1$
 else $\#G_1$
 $\mathcal{L}_D := \mathcal{L}_D \cup \{(c, K)\} \quad \#G_1$

$\mathcal{L}_H := \mathcal{L}_H \cup \{(m, c, K)\}$
return K

Oracle $\text{Decaps}_\perp^{(c^*)}(sk^*, c) \quad \#G_1, G_2$

if $\exists K$ s.t. $(c, K) \in \mathcal{L}_D$
 return K
 $K \leftarrow \mathcal{K}$
 $m' := \text{Dec}_1(sk^*, c)$
if $m' = \perp$
 $K := \perp$
 $\mathcal{L}_D := \mathcal{L}_D \cup \{(c, K)\}$
 $\mathcal{L}_H := \mathcal{L}_H \cup \{(m', c, K)\}$
return K

Figure 2: Sequence of games needed to prove [Theorem 5](#) and the consequential oracle modifications.

Proof. Differently from \mathbf{G}_0 , in game \mathbf{G}_1 we have modified the oracles $\text{Decaps}_{\perp}^{(c^*)}$ and \mathbf{H} in order to avoid the usage of the secret key. In particular, \mathbf{G}_1 defines two sets \mathcal{L}_H and \mathcal{L}_D , where \mathcal{L}_H contains all entries of the form (m, c, K) when either $\text{Decaps}_{\perp}^{(c^*)}$ is queried about a ciphertext c or \mathbf{H} was queried about (m, c) , and \mathcal{L}_D contains all the entries (c, K) when either $\text{Decaps}_{\perp}^{(c^*)}$ is queried about c or \mathbf{H} is queried about (m, c) for $m = \text{Dec}_1(\text{sk}', c)$. Now, we want to show that the view of \mathbf{B} in \mathbf{G}_0 and \mathbf{G}_1 is distributed exactly in the same manner. For this purpose, let us consider a ciphertext c' and a message $m' = \text{Dec}_1(\text{sk}', c')$ for which \mathbf{B} has never been queried $\text{Decaps}_{\perp}^{(c^*)}$:

- **Case $m' = \perp$:** in game \mathbf{G}_0 , $\text{Dec}_1(\text{sk}', c')$ will return \perp to indicate that c' is a malformed ciphertext, which is exactly the behavior the $\text{Decaps}_{\perp}^{(c^*)}$ has in game \mathbf{G}_1 . Regarding the behavior of \mathbf{H} , we can see that in both games \mathbf{H} returns a randomly chosen key.
- **Case $m' \neq \perp$:** in game \mathbf{G}_0 , $\text{Decaps}_{\perp}^{(c^*)}$ returns $K := \mathbf{H}(m', c')$ which is either a fresh key randomly chosen from the key space if (m', c') has never been queried to \mathbf{H} , or taken from \mathcal{L}_H if (m', c') was already stored in \mathcal{L}_H . In game \mathbf{G}_1 we need to consider two sub-cases:
 - \mathbf{B} first queries \mathbf{H} about (m', c') and then queries $\text{Decaps}_{\perp}^{(c^*)}$ about c' : In this case, \mathbf{H} returns a key K which is either a fresh key randomly chosen from the key space, or it is already stored in \mathcal{L}_H . Since $\text{Decaps}_{\perp}^{(c^*)}$ has not been queried about c' yet, \mathbf{H} will add (c, K) to \mathcal{L}_D . Next, when $\text{Decaps}_{\perp}^{(c^*)}$ will be queried about c' , it will return a key K stored in \mathcal{L}_D that coincides with the key stored in \mathcal{L}_H , as in game \mathbf{G}_0 .
 - \mathbf{B} first queries $\text{Decaps}_{\perp}^{(c^*)}$ about c' and then queries \mathbf{H} about (m', c') : In this case, $\text{Decaps}_{\perp}^{(c^*)}$ returns a randomly chosen K , which is added to both \mathcal{L}_D and \mathcal{L}_H . Subsequently, when \mathbf{H} is queried about (m', c') , the oracle will return the key K stored in \mathcal{L}_H that coincides with the key K stored in \mathcal{L}_D . This ensures that $\text{Decaps}_{\perp}^{(c^*)}(c') = \mathbf{H}(m', c') = K$, as in game \mathbf{G}_0 .

Therefore, we have that $\mathbf{G}_1^{\mathbf{B}}(\lambda) \equiv \mathbf{G}_0^{\mathbf{B}}(\lambda)$ □

Lemma 3. $\mathbf{G}_1^{\mathbf{B}}(\lambda) \approx_c \mathbf{G}_2^{\mathbf{B}}(\lambda)$.

Proof. In game \mathbf{G}_2 we make the following two modifications. First, the challenger takes a uniformly sampled key rather than the real key output by the oracle \mathbf{H} , and second we make the oracle \mathbf{H} output \perp when queried about (m^*, c^*) . By denoting the latter event **QUERY**, we notice that \mathbf{G}_1 and \mathbf{G}_2 are identically distributed conditioned to the event **QUERY** not happening. Thus, the only hope for the adversary to distinguish between the \mathbf{G}_1 and \mathbf{G}_2 is to trigger the **QUERY** event in \mathbf{G}_2 . It is easy to see that when \mathbf{B} queries \mathbf{H} in game \mathbf{G}_1 about (m^*, c^*) , \mathbf{H} will return the challenge encapsulation key K^* . Instead, in game \mathbf{G}_2 , the game returns \perp when \mathbf{B} queries \mathbf{H} about (m^*, c^*) . Let us now assume that **QUERY** is not triggered. Since K^* is either an output of \mathbf{H} or a randomly chosen value, the adversary can only try to distinguish by guessing the plaintext m^* of c^* , calculate $K' := \mathbf{H}(m^*, c^*)$ and then check whether K' is equal to K^* . However, this coincides with the **QUERY** event. Alternatively, the adversary might try to distinguish by making \mathbf{G}_1 always output 1, i.e. \mathbf{B} may try to come up with a tuple (pk', R, c') for which the relation $R(K', K^*, \text{pk}', \text{pk}^*, c')$ holds for a key K^* produced by the encapsulation algorithm under pk' (as in \mathbf{G}_1), but does not hold when such key is randomly chosen (as in \mathbf{G}_2). Note that \mathbf{H} is a random oracle, so the set of possible relations is restricted to the ones having $K' = K^* = \mathbf{H}(m^*, c^*)$, otherwise even a single bit different than (m^*, c^*) in the input of \mathbf{H} leads to an independent output. For such a relation the adversary may try to find a public key pk' for which the encapsulation of K^* under pk' leads to c^* . However, for the perfect public-key uniqueness property of Π_1 , such public key does not exist. Since K' can be computed only by giving as input to \mathbf{H} a pair (m', c') in which either $m' \neq m^*$ or $c' \neq c^*$, such a key is independent from the challenge key. Hence, the distribution of the adversary's view in both games is identical given that **QUERY** does not happen. Now, to estimate $\Pr[\text{QUERY}]$ we construct an efficient adversary \mathbf{A} breaking OW-PCVA of Π_1 when **QUERY** occurs. In particular, we define \mathbf{A} in Fig. 3. Notice that \mathbf{A} perfectly simulates \mathbf{G}_1 . Indeed, the occurrence

<p>Adversary $A^{\text{PCO}(\text{sk}^*, \cdot, \cdot), \text{CVO}^{(c^*)}(\text{sk}^*, \cdot)}(\text{pk}, c^*)$</p> <hr style="border: 0.5px solid black;"/> <p>$K^* \leftarrow \mathcal{K}$</p> <p>$(\text{pk}', R, c') \leftarrow \mathcal{B}_2^{\text{Decaps}_{\perp}^{(c^*)}(\cdot), \text{H}(\cdot, \cdot)}(\text{pk}^*, c^*, K^*)$</p> <p>if $\exists (m, c^*) \in \mathcal{L}_H$ s.t. $\text{PCO}(\text{sk}^*, m, c^*) = 1$</p> <p style="padding-left: 20px;">return m</p> <p>return \perp</p> <p>Oracle $\text{Decaps}_{\perp}^{(c^*)}(\text{sk}, c)$</p> <hr style="border: 0.5px solid black;"/> <p>if $\exists c$ s.t. $(c, K) \in \mathcal{L}_D$</p> <p style="padding-left: 20px;">return K</p> <p>if $\text{CVO}^{(c^*)}(\text{sk}^*, c) = 0$</p> <p style="padding-left: 20px;">return \perp</p> <p>$K \leftarrow \mathcal{K}$</p> <p>$\mathcal{L}_D := \mathcal{L}_D \cup \{(c, K)\}$</p> <p>return K</p>	<p>Oracle $\text{H}(m, c)$</p> <hr style="border: 0.5px solid black;"/> <p>if $\exists K$ s.t. $(m, c, K) \in \mathcal{L}_H$</p> <p style="padding-left: 20px;">return K</p> <p>$K \leftarrow \mathcal{K}$</p> <p>if $\text{PCO}(\text{sk}^*, m, c) = 1$</p> <p style="padding-left: 20px;">if $\exists K'$ s.t. $(c, K') \in \mathcal{L}_D$</p> <p style="padding-left: 40px;">$K := K'$</p> <p style="padding-left: 20px;">else</p> <p style="padding-left: 40px;">$\mathcal{L}_D := \mathcal{L}_D \cup \{(c, K)\}$</p> <p style="padding-left: 20px;">$\mathcal{L}_H := \mathcal{L}_H \cup \{(m, c)\}$</p> <p style="padding-left: 20px;">return K</p>
--	--

Figure 3: Adversary A breaking security of the underlying OW-PCVA PKE.

of QUERY implies that B has queried H about (m, c) , in which $(m, c) \in \mathcal{L}_H$ for $m = m^*$ and $c = c^*$. A will return $m = m^*$ and win the OW-PCVA experiment. Since such condition coincides with the QUERY event, we get that the probability of B of triggering QUERY coincides with the probability of A in winning the OW-PCVA experiment, i.e.

$$\Pr[\text{QUERY}] = \Pr\left[\text{PKE}_{\Pi_1, A}^{\text{ow-pcva}}(\lambda)\right].$$

□

Lemma 4. $\mathbf{G}_2^{\mathcal{B}}(\lambda) \approx_c \mathbf{G}_3^{\mathcal{B}}(\lambda)$.

Proof. Notice that game \mathbf{G}_3 is identically distributed to \mathbf{G}_0 , except that the challenge encapsulation key is randomly chosen from the key space. We will show that the view of B in \mathbf{G}_2 and \mathbf{G}_3 is identically distributed, under the condition that the QUERY event does not occur.

Let us fix a ciphertext c' and a message $m' = \text{Dec}_2(\text{sk}^*, c')$. We consider two cases:

- **Case** $m' \notin \mathcal{M}$: in game \mathbf{G}_3 , when $\text{Decaps}_{\perp}^{(c^*)}$ is queried about c' , it will return \perp , which is exactly what $\text{Decaps}_{\perp}^{(c^*)}$ returns in \mathbf{G}_2 .
- **Case** $m' \in \mathcal{M}$. Here, we need to consider two sub-cases:
 - B first queries $\text{Decaps}_{\perp}^{(c^*)}$ and then H: assume that neither $\text{Decaps}_{\perp}^{(c^*)}$ nor H have been queried before about c' and (m', c') respectively. In \mathbf{G}_2 , when B queries $\text{Decaps}_{\perp}^{(c^*)}$, $\text{Decaps}_{\perp}^{(c^*)}$ will return a uniform key K , add an entry of the form (c', K) to \mathcal{L}_D and an entry of the form (m', c', K) to \mathcal{L}_H . Next, when B queries H about (m', c', pk') , H will return the same key returned by $\text{Decaps}_{\perp}^{(c^*)}$ since $(m', c', K) \in \mathcal{L}_H$. In \mathbf{G}_3 , $\text{Decaps}_{\perp}^{(c^*)}$ will return a uniform key $K = \text{H}(m', c')$. Next, when B queries H about (m', c') , H will return the same key K . Therefore, in both games we have $\text{Decaps}_{\perp}^{(c^*)} = K = \text{H}(m', c')$.
 - B first queries H and then $\text{Decaps}_{\perp}^{(c^*)}$: As before, let us assume that neither $\text{Decaps}_{\perp}^{(c^*)}$ nor H have been queried before about c' and (m', c') respectively. In \mathbf{G}_2 , when B queries H, H will return

a uniform key K and add an entry of the form (c', K) to \mathcal{L}_D . Next, when $\text{Decaps}_{\perp}^{(c^*)}$ is queried about c' , $\text{Decaps}_{\perp}^{(c^*)}$ will return the same key returned by H , since $(c', K) \in \mathcal{L}_D$. In game \mathbf{G}_3 when B queries H , H will return a uniform key K and add an entry of the form (m', c', K) to \mathcal{L}_H . Next, when $\text{Decaps}_{\perp}^{(c^*)}$ is queried about c' , $\text{Decaps}_{\perp}^{(c^*)}$ will return the same key K , due to the fact that $(m', c', K) \in \mathcal{L}_H$. Thus, in both games we have $H(m', c') = K = \text{Decaps}_{\perp}^{(c^*)}$.

Since the only hope for the adversary is to trigger the QUERY event in \mathbf{G}_3 , the probability for B to distinguish between \mathbf{G}_2 and \mathbf{G}_3 is bounded by the probability of winning the OW-PCVA game of the underlying PKE scheme. \square

Lemma 5. $\mathbf{G}_3^B(\lambda) \equiv \text{KEM}_{\tilde{\Pi}_1, B, \mathcal{S}}^{nm-cca2^*}(\lambda)$.

Proof. Since \mathbf{G}_3 is similar to \mathbf{G}_0 , with the only difference that the encapsulation key K^* is uniform and independent from the one obtained by querying H , for the same considerations that we did for \mathbf{G}_0 , it holds that the two distributions are identically distributed. \square

Combining the above lemmas, we get that $\text{KEM}_{\tilde{\Pi}_1^{\perp}, B}^{nm-cca2^*}(\lambda) \equiv \mathbf{G}_0^B(\lambda) \equiv \mathbf{G}_1^B(\lambda) \approx_c \mathbf{G}_2^B(\lambda) \approx_c \mathbf{G}_3^B(\lambda) \equiv \text{KEM}_{\tilde{\Pi}_1^{\perp}, B, \mathcal{S}}^{nm-cca2^*}(\lambda)$, thus

$$\text{KEM}_{\tilde{\Pi}_1^{\perp}, B}^{nm-cca2^*}(\lambda) \approx_c \text{KEM}_{\tilde{\Pi}_1^{\perp}, B, \mathcal{S}}^{nm-cca2^*}(\lambda).$$

\square

4.3 Modified transformation \hat{U}^{\perp}

In the following, we leverage the idea of prefix hashing introduced by Duman et al. [10], to construct a completely non-malleable KEM without requiring public-key uniqueness of the underlying PKE. In particular, our \hat{U}^{\perp} is identical to U^{\perp} , except that now the encapsulation algorithm gives as input to the random oracle H also the public key pk together with the message m and the ciphertext c from the underlying PKE scheme, i.e. $K := H(m, c, pk)$. Note that now the decapsulation oracle must take as input also the challenge public key together with the challenge secret key in order to recompute $H(m, c, pk)$. The theorem below states that $\tilde{\Pi}_1^{\perp} := \hat{U}^{\perp}[\Pi_1, H]$ is completely non-malleable. Since the techniques used to prove the theorem below are similar to the ones used to prove [Theorem 5](#), we will only highlight the changes in the sequence of games w.r.t. the proof of [Theorem 5](#), and the changes needed to prove some lemma when required.

Theorem 6 (Π_1 OW-PCVA $\xrightarrow{ROM} \tilde{\Pi}_1^{\perp}$ NM-CCA2*). *Assuming the existence of a random oracle H , if Π_1 is a correct OW-PCVA secure PKE ([Definition 3](#)), then $\tilde{\Pi}_1^{\perp}$ defined above is a correct NM-ATK* secure KEM ([Definition 8](#)).*

Proof. The sequence of games and the consequential differences in the oracles are described in [Fig. 4](#).

Lemma 6. $\text{KEM}_{\tilde{\Pi}_1^{\perp}, B}^{nm-cca2^*}(\lambda) \equiv \mathbf{G}_0^B(\lambda)$.

Proof. The proof is identical to the one of [Lemma 1](#). \square

Lemma 7. $\mathbf{G}_0^B(\lambda) \equiv \mathbf{G}_1^B(\lambda)$.

Proof. The proof is identical to the one of [Lemma 2](#). \square

Lemma 8. $\mathbf{G}_1^B(\lambda) \approx_c \mathbf{G}_2^B(\lambda)$.

Games $\mathbf{G}_0^{\mathbf{B}}(\lambda) - \mathbf{G}_3^{\mathbf{B}}(\lambda)$

$(pk^*, sk^*) \leftarrow \mathcal{G}_{\text{Gen}_1}(1^\lambda)$
 $st \leftarrow \mathcal{B}_1^{\text{Decaps}_{\perp}^{(c^*)}(sk^*, \cdot), H(\cdot, \cdot)}(pk^*)$
 $m^* \leftarrow \mathcal{M}$
 $c^* \leftarrow \text{Enc}_1(pk, m^*)$
 $K^* := H(m^*, c^*, pk^*) \quad \boxed{\#G_0, G_1}$
 $K^* \leftarrow \{0, 1\}^\lambda \quad \boxed{\#G_2, G_3}$
 $(pk', R, c') \leftarrow \mathcal{B}_2^{\text{Decaps}_{\perp}^{(c^*)}(sk^*, \cdot), H(\cdot, \cdot)}(pk^*, c^*, K^*, st)$
return 1 iff $\exists(K', r)$ such that
 $(c', K') := \text{Encaps}(pk'; r) \wedge (K' \neq \perp)$
 $\wedge (pk' \neq pk^* \vee K' \neq K^*)$
 $\wedge R(K', K^*, pk', pk^*)$

Oracle $\text{Decaps}_{\perp}^{(c^*)}(sk^*, pk^*, c) \quad \boxed{\#G_0, G_3}$

$m := \text{Dec}_1(sk^*, c)$
if $m' = \perp$
 return \perp
return $K := H(m', c, pk^*)$

Oracle $H(m, c, pk)$

if $\exists K$ s.t. $(m, c, pk, K) \in \mathcal{L}_H$
 return K
 $K \leftarrow \mathcal{K}$
if $\text{Dec}_1(sk^*, c) = m \quad \boxed{\#G_1, G_2}$
 if $c = c^*$
 return $\perp \quad \boxed{\#G_2}$
 if $\exists K'$ s.t. $(c, K') \in \mathcal{L}_D \quad \boxed{\#G_1}$
 $K := K' \quad \boxed{\#G_1}$
 else $\boxed{\#G_1}$
 $\mathcal{L}_D := \mathcal{L}_D \cup \{(c, K)\} \quad \boxed{\#G_1}$

$\mathcal{L}_H := \mathcal{L}_H \cup \{(m, c, pk, K)\}$

return K

Oracle $\text{Decaps}_{\perp}^{(c^*)}(sk^*, pk^*, c) \quad \boxed{\#G_1, G_2}$

if $\exists K$ s.t. $(c, K) \in \mathcal{L}_D$
 return K
 $K \leftarrow \mathcal{K}$
 $m' := \text{Dec}_1(sk^*, c)$
if $m' = \perp$
 $K := \perp$
 $\mathcal{L}_D := \mathcal{L}_D \cup \{(c, K)\}$
 $\mathcal{L}_H := \mathcal{L}_H \cup \{(m', c, pk^*, K)\}$
return K

Figure 4: Sequence of games needed to prove [Theorem 6](#) and the consequential modifications of the oracles.

Proof. The proof is identical to [Lemma 3](#), except that now we do not use public key uniqueness of the underlying PKE scheme to argue that the \mathbf{G}_1 and \mathbf{G}_2 are identically distributed conditioned to the fact that the event QUERY does not happen. In this case, the adversary can try to distinguish between \mathbf{G}_1 and \mathbf{G}_2 by guessing the plaintext m^* of c^* , calculate $K' := H(m^*, c^*, \text{pk}^*)$ and then check whether K' is equal to K^* or not. However, this coincides with the QUERY event. Alternatively, the adversary might try to distinguish making \mathbf{G}_1 always output 1, i.e. B tries to come up with a tuple (pk', R, c') for which the relation $R(K', K^*, \text{pk}', \text{pk}^*, c')$ holds for a key K^* encapsulated by c^* under pk' (as in \mathbf{G}_1), but does not hold when such key is randomly chosen (as in \mathbf{G}_2). However, for the random oracle assumption, since pk' is part of the input of H , the key $K' := H(m^*, c^*, \text{pk}')$ will be independent from $K^* := H(m^*, c^*, \text{pk}^*)$. Thus, the distribution of \mathbf{G}_1 and \mathbf{G}_2 is identical when QUERY is not triggered.

As in [Lemma 3](#), to estimate $\Pr[\text{QUERY}]$ we construct an efficient adversary A breaking OW-PCVA of PKE_1 when QUERY occurs. We define A in [Fig. 5](#). Notice that A perfectly simulates \mathbf{G}_1 . Indeed, the

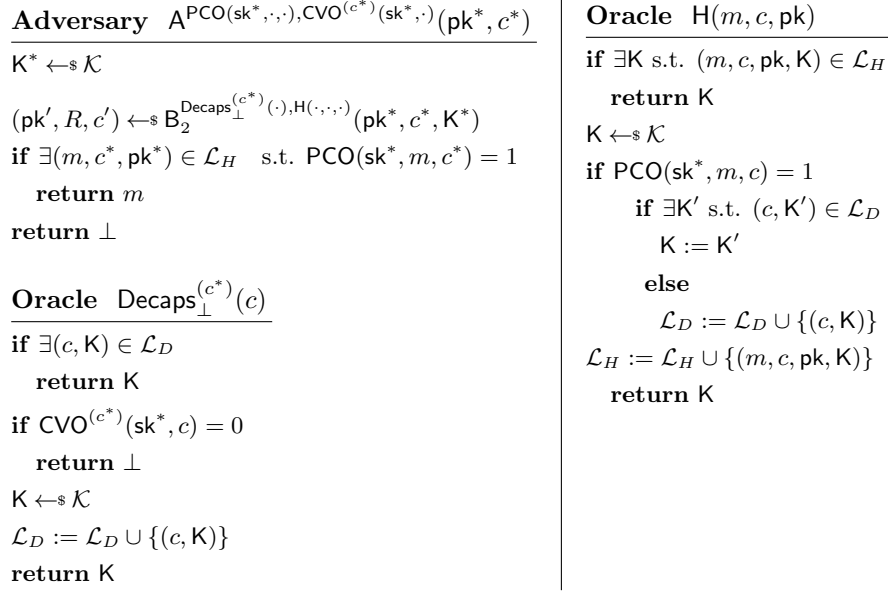


Figure 5: Adversary A breaking security of the underlying OW-PCVA PKE.

occurrence of QUERY implies that B has queried H about (m, c, pk) , in which $(m, c, \text{pk}) \in \mathcal{L}_H$ for $m = m^*$, $c = c^*$ and $\text{pk} = \text{pk}^*$. A then returns $m = m^*$. Since such event coincides with QUERY, we get that the probability of B of triggering QUERY coincides with the probability of A in winning the OW-PCVA experiment, i.e.

$$\Pr[\text{QUERY}] = \Pr\left[\text{PKE}_{\Pi_1, A}^{\text{ow-pcva}}(\lambda)\right] \leq \text{negl}(\lambda).$$

□

Lemma 9. $\mathbf{G}_2^B(\lambda) \approx_c \mathbf{G}_3^B(\lambda)$.

Proof. The proof is identical to the one of [Lemma 4](#).

□

Lemma 10. $\mathbf{G}_3^B(\lambda) \equiv \text{KEM}_{\Pi_1, B, \S}^{nm-cca2^*}(\lambda)$.

Proof. Since \mathbf{G}_3 is similar to \mathbf{G}_0 , with the only difference that the encapsulation key K^* is uniform it is independent from the one obtained by querying H , for the same considerations that we did for \mathbf{G}_0 , it holds that the two distributions are identically distributed.

□

Combining the above lemmas, we get that $\text{KEM}_{\Pi_1^\perp, \mathcal{B}}^{nm-cca2^*}(\lambda) \equiv \mathbf{G}_0^{\mathcal{B}}(\lambda) \equiv \mathbf{G}_1^{\mathcal{B}}(\lambda) \approx_c \mathbf{G}_2^{\mathcal{B}}(\lambda) \approx_c \mathbf{G}_3^{\mathcal{B}}(\lambda) \equiv \text{KEM}_{\Pi_1^\perp, \mathcal{B}, \mathcal{S}}^{nm-cca2^*}(\lambda)$, thus

$$\text{KEM}_{\Pi_1^\perp, \mathcal{B}}^{nm-cca2^*}(\lambda) \approx_c \text{KEM}_{\Pi_1^\perp, \mathcal{B}, \mathcal{S}}^{nm-cca2^*}(\lambda).$$

□

5 Relation with Completely Non-Malleable PKE

In the following, we will show which kind of relationships exist between a NM-ATK* PKE schemes and NM-ATK* KEM schemes. We will proceed by first showing that a NM-ATK* PKE scheme is intrinsically a NM-ATK* KEM scheme, and next we will introduce a construction to prove that a NM-ATK* KEM can be used together with a NM-ATK SKE scheme to construct a NM-ATK* PKE scheme by using the KEM/DEM paradigm.

5.1 NM-ATK* PKE \implies NM-ATK* KEM

Wlog, we can assume that $\mathcal{M} = \mathcal{K}$. Let $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ be a PKE scheme, and $\Pi' = (\text{Gen}, \text{Encaps}, \text{Decaps})$ be a KEM scheme defined as follows:

<p style="text-align: center; margin: 0;"><u>Algorithm Encaps(pk)</u></p> <p style="margin: 0;">$K \leftarrow_s \mathcal{K}$</p> <p style="margin: 0;">$c \leftarrow_s \text{Enc}(\text{pk}, K)$</p> <p style="margin: 0;">return (c, K)</p>	<p style="text-align: center; margin: 0;"><u>Algorithm Decaps(sk, c)</u></p> <p style="margin: 0;">$K := \text{Dec}(\text{sk}, c)$</p> <p style="margin: 0;">return K</p>
---	---

Theorem 7 (NM-ATK* PKE \implies NM-ATK* KEM). *If Π is a NM-ATK* secure PKE (Definition 4) with respect to a set of relations \mathcal{R} , then Π' is a NM-ATK* secure KEM (Definition 8) with respect to \mathcal{R} .*

Proof. Let us assume that there exists a PPT adversary A that breaks NM-ATK* security of Π' with non-negligible probability, then we can build an efficient distinguisher D that breaks NM-ATK* security of Π . The intuition behind the proof is that D will choose a distribution over the message space in such a way that only two messages, say K and K' , can be sampled by the the challenger playing NM-ATK* of the PKE scheme Π . In particular, D does the following:

1. Takes as input a public key pk and chooses a message distribution \mathcal{M} from which only two messages can be chosen, which we call K and K' .
2. Takes as input a challenge ciphertext c which is either an encryption of K under pk or an encryption of K' under pk .
3. Run $A(\text{pk}, c, K)$. When A asks a decapsulation-oracle query for a ciphertext \hat{c} do the following:
 - (a) Query the decryption oracle $\text{Dec}(\text{sk}^*, \cdot)$ about \hat{c} to obtain a key \hat{K} .
 - (b) Return \hat{K} to A .
4. When A returns (pk', R, c') , return (pk', R, c') as well.

Let us analyze the behavior of D . First, notice that, independently from c , D will always send the tuple (pk, c, K) to A . Recall also that the message space of c is restricted to messages K and K' . In particular:

- when c is an encryption of K under pk , D the view of A when run as a subroutine of D is identically distributed to its view in $\text{KEM}_{\Pi', A}^{nm-atk^*}(\lambda)$.

- When c is an encryption of K' under pk , since K' is uniform and independent from K , the view of A when run as a subroutine of D is distributed identically to its view in $\text{KEM}_{\Pi', A, \$}^{nm-atk^*}(\lambda)$.

To summarize, the probability that D distinguishes between $\text{PKE}_{\Pi, D}^{nm-atk^*}(\lambda)$ and $\text{PKE}_{\Pi, D, \$}^{nm-atk^*}(\lambda)$ is the same of A distinguishing between $\text{KEM}_{\Pi', A, \$}^{nm-atk^*}(\lambda)$ and $\text{KEM}_{\Pi', A, \$}^{nm-atk^*}(\lambda)$, that we assumed to be non-negligible. This leads to a contradiction due to the NM-ATK* security of Π . \square

5.2 NM-ATK* KEM + NM-ATK SKE \implies NM-ATK* PKE

It is well known that a KEM scheme alone doesn't allow to build a PKE scheme, due to the fact that an encapsulation algorithm can be instantiated by encrypting a uniformly chosen message. To solve this issue, a secret-key encryption scheme can be used along with a key-encapsulation mechanism.

Let $\Pi^{kem} = (\text{Gen}^{kem}, \text{Encaps}, \text{Decaps})$ be a NM-ATK* KEM with key space \mathcal{K} , and $\Pi^{ske} = (\text{Gen}^{ske}, \text{Enc}, \text{Dec})$ be a NM-ATK SKE scheme with message space \mathcal{M} and the same key space \mathcal{K} of Π^{kem} , we can construct a hybrid PKE scheme $\Pi^{hy} = (\text{Gen}^{hy}, \text{Enc}^{hy}, \text{Dec}^{hy})$ as defined below.

Algorithm $\text{Gen}^{hy}(1^\lambda)$	Algorithm $\text{Enc}^{hy}(\text{pk}, m)$	Algorithm $\text{Dec}^{hy}(\text{sk}, c)$
$(\text{pk}, \text{sk}) \leftarrow_s \text{Gen}^{kem}(1^\lambda)$	$(c, K) \leftarrow_s \text{Encaps}(\text{pk})$	$K := \text{Decaps}(\text{sk}, c)$
return (pk, sk)	$c' \leftarrow_s \text{Enc}(K, m)$	return $m := \text{Dec}(K, c)$
	return (c, c')	

We prove that if Π^{kem} is a completely non-malleable KEM and Π^{ske} is a non-malleable SKE, then Π^{hy} is a completely non-malleable PKE.

Theorem 8. *If Π^{kem} is an NM-ATK* secure KEM (Definition 8) with respect to a set of relations \mathcal{R} and Π^{ske} is an NM-ATK secure SKE (Definition 6) with respect to a set of relations $\mathcal{R}' \subseteq \mathcal{R}$, then the scheme Π^{hy} described above is a NM-ATK* secure PKE scheme (Definition 4) with respect to \mathcal{R} .*

Proof. Correctness of the obtained PKE follows from the ϵ -correctness of the underlying KEM and SKE schemes. The idea behind the proof is that, given the challenge ciphertext $c^* = (c_1^*, c_2^*)$, we can use NM-ATK* security of Π^{kem} to decouple the key encapsulated in c_1^* from the key used in c_2^* to encrypt the message with the underlying SKE scheme Π^{ske} . At this point, since the encapsulated key is randomly chosen and independent from the encryption key, it is not possible for A to distinguish between a correct encryption (i.e., where the encapsulated key and the encryption key are the same) and an encryption where the key encapsulated in c^* is randomly chosen and independent from the one used to encrypt c_2^* . This holds even if A is allowed to maul pk^* into some related public key pk' . The next step is to use the NM-ATK security of Π^{ske} to decouple m^* from the relation R , i.e. given a ciphertext c_2^* encrypting m^* , it is infeasible for an adversary to distinguish between the experiment where the relation R was checked by using either m^* or \tilde{m} . Finally, NM-ATK* security of Π^{kem} can be used to re-join together the key encapsulated in c_1^* with the key used to encrypt m in c_2^* . Let $A = A^{hy}$, the sequence of games is described in Fig. 6. The part of the proof required for a specific flavour of NM-ATK* will be highlighted with a tag [NM-ATK*].

Lemma 11. $\mathbf{G}_0(\lambda) \approx_c \mathbf{G}_1(\lambda)$.

Proof. Let us assume that A^{hy} can distinguish between \mathbf{G}_0 and \mathbf{G}_1 with non-negligible probability. We can construct an adversary A^{nm} breaking NM-ATK* security of Π^{kem} . A^{nm} behaves as follows:

1. Take as input a public key pk^* , a ciphertext c and a key \hat{K} , where either $\hat{K} = K^*$ (the key encapsulated in c) or $\hat{K} \leftarrow_s \mathcal{K}$.

Game $G_0^A(\lambda)$

$(pk^*, sk^*) \leftarrow_s \text{Gen}(1^\lambda)$
 $(\mathcal{M}, s) \leftarrow_s A_1^{O_1}(\text{pk})$
 $m^* \leftarrow_s \mathcal{M}$
 $(c_1^*, K^*) \leftarrow_s \text{Encaps}(pk^*)$
 $c_2^* \leftarrow_s \text{Enc}(K^*, m^*)$
 $(pk, R, c) \leftarrow_s A_2^{O_2}(\mathcal{M}, pk^*, s, c^*)$
return 1 iff $\exists(m, r)$ s.t.
 $(c = \text{Enc}^{hy}(pk, m; r)) \wedge$
 $(c \neq c^* \vee pk' \neq pk) \wedge$
 $(m \neq \perp) \wedge R(m, m^*, pk, pk^*, c)$

Game $G_2^A(\lambda)$

$(pk^*, sk^*) \leftarrow_s \text{Gen}(1^\lambda)$
 $(\mathcal{M}, s) \leftarrow_s A_1^{O_1}(\text{pk})$
 $m^*, \tilde{m} \leftarrow_s \mathcal{M}$
 $\hat{K} \leftarrow_s \mathcal{K}$
 $(c_1^*, K^*) \leftarrow_s \text{Encaps}(pk^*)$
 $c_2^* \leftarrow_s \text{Enc}(\hat{K}, m^*)$
 $(pk, R, c) \leftarrow_s A_2^{O_2}(\mathcal{M}, pk^*, s, c^*)$
return 1 iff $\exists(m, r)$ s.t.
 $(c = \text{Enc}^{hy}(pk, m; r)) \wedge$
 $(c \neq c^* \vee pk' \neq pk) \wedge$
 $(m \neq \perp) \wedge R(m, \tilde{m}, pk, pk^*, c)$

Game $G_1^A(\lambda)$

$(pk^*, sk^*) \leftarrow_s \text{Gen}(1^\lambda)$
 $(\mathcal{M}, s) \leftarrow_s A_1^{O_1^{(\cdot)}}(\text{pk})$
 $m^* \leftarrow_s \mathcal{M}$
 $\hat{K} \leftarrow_s \mathcal{K}$
 $(c_1^*, K^*) \leftarrow_s \text{Encaps}(pk^*)$
 $c_2^* \leftarrow_s \text{Enc}(\hat{K}, m^*)$
 $(pk, R, c) \leftarrow_s A_2^{O_2^{(\cdot)}}(\mathcal{M}, pk^*, s, c^*)$
return 1 iff $\exists(m, r)$ s.t.
 $(c = \text{Enc}^{hy}(pk, m; r)) \wedge$
 $(c \neq c^* \vee pk' \neq pk) \wedge$
 $(m \neq \perp) \wedge R(m, m^*, pk, pk^*, c)$

Game $G_3^A(\lambda)$

$(pk^*, sk^*) \leftarrow_s \text{Gen}(1^\lambda)$
 $(\mathcal{M}, s) \leftarrow_s A_1^{O_1^{(\cdot)}}(\text{pk})$
 $m^*, \tilde{m} \leftarrow_s \mathcal{M}$
 $(c_1^*, K^*) \leftarrow_s \text{Encaps}(pk^*)$
 $c_2^* \leftarrow_s \text{Enc}(K^*, m^*)$
 $(pk, R, c) \leftarrow_s A_2^{O_2^{(\cdot)}}(\mathcal{M}, pk^*, s, c^*)$
return 1 iff $\exists(m, r)$ s.t.
 $(c = \text{Enc}^{hy}(pk, m; r)) \wedge$
 $(c \neq c^* \vee pk' \neq pk) \wedge$
 $(m \neq \perp) \wedge R(m, \tilde{m}, pk, pk^*, c)$

Figure 6: Sequence of games needed to prove [Theorem 8](#).

2. Run $A^{hy}(\text{pk}^*)$.
 [NM-CCA*1/NM-CCA2*] When A^{hy} asks a decryption-oracle query about a ciphertext (c_1, c_2) query $\text{Decaps}(\text{sk}^*, \cdot)$ about c_1 to obtain a key K' and return $m := \text{Dec}(K', c_2)$.
 When A^{hy} outputs a message distribution \mathcal{M} , take a uniform message $m \leftarrow_s \mathcal{M}$, compute $c \leftarrow_s \text{Enc}(K, m)$ and return $c^* = (c, c')$ to A^{hy} .
3. [NM-CCA2*] When A^{hy} asks a decryption-oracle query about a ciphertext (c_1, c_2) do the following:
 - if $c_1 = c$ and $c_2 = c'$ then return $m := \perp$ (i.e the query is not admissible).
 - if $c_1 = c$ and $c_2 \neq c'$ then return $m := \text{Dec}(\hat{K}, c_2)$.
 - else, query $\text{Decaps}^{(c^*)}(\text{sk}^*, \cdot)$ about c_1 to obtain a key K' and return $m := \text{Dec}(K', c_2)$.
4. When A^{hy} outputs $(\text{pk}, R, (c_1, c_2))$, output (pk, R', c_1) , where $R'(\cdot, \cdot, \text{pk}, \text{pk}^*, c) = R(\cdot, \cdot, \text{pk}, \text{pk}^*, c)$.

Notice that, since the only difference between \mathbf{G}_0 and \mathbf{G}_1 is that in \mathbf{G}_1 the key is chosen at random, the only hope for the adversary A^{hyb} to distinguish between the two hybrids is by finding a relation holding between pk , pk^* , and c that is satisfied in \mathbf{G}_0 but not in \mathbf{G}_1 (or vice-versa). Hence, $R'(\cdot, \cdot, \text{pk}, \text{pk}^*, c) = R(\cdot, \cdot, \text{pk}, \text{pk}^*, c)$ is indeed a suitable relation for A^{nm} . When \hat{K} taken as input by A^{nm} is K^* , then A^{nm} perfectly simulates \mathbf{G}_0 . When \hat{K} taken as input by A^{nm} is randomly chosen, A^{nm} perfectly simulates \mathbf{G}_1 . If A^{hy} distinguishes between \mathbf{G}_0 and \mathbf{G}_1 with non-negligible probability, then A^{nm} breaks NM-ATK* security of the underlying KEM scheme with non-negligible probability. This leads to a contradiction. \square

Lemma 12. $\mathbf{G}_1(\lambda) \approx_c \mathbf{G}_2(\lambda)$.

Proof. Let us assume that A^{hy} can distinguish between \mathbf{G}_1 and \mathbf{G}_2 with non-negligible probability, we construct an adversary A^{atk} breaking NM-ATK security of Π^{ske} . A^{atk} behaves as follows:

1. Receive as input a key K .
2. Generate a pair $(\text{pk}^*, \text{sk}^*) \leftarrow_s \text{Gen}^{kem}(1^\lambda)$.
3. Run $A^{hy}(\text{pk}^*)$.
 [NM-CCA*1/NM-CCA2*] When A^{hy} asks a decryption-oracle query about a ciphertext (c_1, c_2) , query the decryption oracle $\text{Dec}(\hat{K}, \cdot)$ of the NM-CCA experiment about c_2 to obtain a message m .
4. When A^{hy} outputs a message distribution \mathcal{M} , output \mathcal{M} to the challenger.
5. When receiving a ciphertext c' from the challenger, compute $(c_1^*, K^*) \leftarrow_s \text{Encaps}(\text{pk}^*)$, and output (c_1^*, c') to A^{hy} .
6. [NM-CCA2*] When A^{hy} asks a decryption-oracle query about a ciphertext (c_1, c_2) do the following:
 - if $c_1 = c_1^*$ and $c_2 = c'$ then return $m := \perp$ (i.e the query is not admissible).
 - else, query $\text{Dec}^{(c')}(K, \cdot)$ about c_2 to obtain $m := \text{Dec}(K', c_2)$. Then, outputs m to A^{hyb} .
7. When A^{hy} outputs $(\text{pk}, R, (c_1, c_2))$, A^{atk} output (R', c_2) to the challenger, where $R'(m_0, m_1) = R(m_0, m_1, \cdot, \cdot, \cdot)$. The challenger either checks $R'(m, m^*) = 1$ where $m^* := \text{Dec}(\hat{K}, c')$ if A^{atk} is in $\text{SKE}_{\Pi^{ske}, A^{atk}}^{nm-atk}$ or checks if $R'(m, \tilde{m}) = 1$ if \tilde{m} is a randomly chosen message independent from c^* . Note that the only difference between \mathbf{G}_1 and \mathbf{G}_2 is that the game checks that m^* given as an input to R is encrypted in c_2^* , whereas in \mathbf{G}_2 the relation R takes as an input a random \tilde{m} . Thus, the only hope for A^{hyb} in distinguishing between the two hybrids is by finding a relation holding between m and m^* but not between m and \tilde{m} (or vice-versa). Thus, we are allowed to cast $R'(m_0, m_1)$ as $R(m_0, m_1, \cdot, \cdot, \cdot)$. When the relation R takes m^* in input, then A^{atk} perfectly simulates \mathbf{G}_1 . When the relation R takes a random \tilde{m} in input, A^{atk} perfectly simulates \mathbf{G}_2 .

If A^{hy} distinguishes between \mathbf{G}_1 and \mathbf{G}_2 with non-negligible probability, then A^{atk} breaks NM-ATK security of the underlying SKE scheme with non-negligible probability. This leads to a contradiction. \square

Lemma 13. $\mathbf{G}_2^A(\lambda) \approx_c \mathbf{G}_3^A(\lambda)$

Proof. Let us assume that A^{hy} can distinguish between \mathbf{G}_2 and \mathbf{G}_3 with non-negligible probability, we construct an adversary A^{nm} breaking NM-ATK* security of Π^{kem} . A^{nm} behaves as follows:

1. Takes as input a public key \mathbf{pk}^* , a ciphertext c and a key \hat{K} , where either $\hat{K} = K^*$ (the key encapsulated in c) or $\hat{K} \leftarrow \mathcal{K}$.
2. Run $A^{hy}(\mathbf{pk}^*)$.
[NM-CCA*1/NM-CCA2*] When A^{hy} asks a decryption-oracle query about a ciphertext (c_1, c_2) query $\text{Decaps}(\mathbf{sk}^*, \cdot)$ about c_1 to obtain a key K' and return $m := \text{Dec}(K', c_2)$.
3. When A^{hy} outputs a message distribution \mathcal{M} , take two uniform messages $m, \tilde{m} \leftarrow \mathcal{M}$, compute $c \leftarrow \text{Enc}(K, m)$ and return $c^* = (c, c')$ to A^{hy} .
4. [NM-CCA2*] When A^{hy} asks a decryption-oracle query about a ciphertext (c_1, c_2) do the following:
 - if $c_1 = c$ and $c_2 = c'$ then return $m := \perp$ (i.e the query is not admissible).
 - if $c_1 = c$ and $c_2 \neq c'$ then return $m := \text{Dec}(\hat{K}, c_2)$.
 - else, query $\text{Decaps}^{(c^*)}(\mathbf{sk}^*, \cdot)$ about c_1 to obtain a key K' and return $m := \text{Dec}(K', c_2)$.
5. When A^{hy} outputs $(\mathbf{pk}, R, (c_1, c_2))$, output (\mathbf{pk}, R', c_1) , where $R'(\cdot, \cdot, \mathbf{pk}, \mathbf{pk}^*, c) = R(\cdot, \cdot, \mathbf{pk}, \mathbf{pk}^*, c)$.

Notice that, since the only difference between \mathbf{G}_2 and \mathbf{G}_3 is that in \mathbf{G}_2 the key is chosen at random, the only hope for the adversary A^{hyb} to distinguish between the two hybrids is by finding a relation holding between $\mathbf{pk}, \mathbf{pk}^*$, and c that is satisfied in \mathbf{G}_2 but not in \mathbf{G}_3 (or vice-versa). Hence, R' is a suitable relation for A^{nm} . When the key \hat{K} taken as input by A^{nm} is randomly chosen, it perfectly simulates \mathbf{G}_2 . When the key \hat{K} taken as input by A^{nm} is K^* , then A^{nm} perfectly simulates \mathbf{G}_3 . If A^{hy} distinguishes between \mathbf{G}_2 and \mathbf{G}_3 with non-negligible probability, then A^{nm} breaks NM-ATK* security of the underlying KEM scheme with non-negligible probability. This leads to a contradiction. \square

It is easy to see that $\mathbf{G}_0^A(\lambda) \equiv \text{PKE}_{\Pi^{hy}, A}^{nm-atk^*}(\lambda)$ and that $\mathbf{G}_3^A(\lambda) \equiv \text{PKE}_{\Pi^{hy}, A, \mathcal{S}}^{nm-atk^*}(\lambda)$. by combining the above lemmas, $\text{PKE}_{\Pi^{hy}, A}^{nm-atk^*}(\lambda) \equiv \mathbf{G}_0^A(\lambda) \approx_c \mathbf{G}_1^A(\lambda) \approx_c \mathbf{G}_2^A(\lambda) \approx_c \mathbf{G}_3^A(\lambda) \equiv \text{PKE}_{\Pi^{hy}, A, \mathcal{S}}^{nm-atk^*}(\lambda)$. Hence, $\text{PKE}_{\Pi^{hy}, A}^{nm-atk^*}(\lambda) \approx_c \text{PKE}_{\Pi^{hy}, A, \mathcal{S}}^{nm-atk^*}(\lambda)$. \square

References

- [1] Barbosa, M., Farshim, P.: Relations among notions of complete non-malleability: Indistinguishability characterisation and efficient construction without random oracles. In: Steinfeld, R., Hawkes, P. (eds.) ACISP (2010)
- [2] Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption. In: 38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997. pp. 394–403. IEEE Computer Society (1997)
- [3] Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations among notions of security for public-key encryption schemes. In: Krawczyk, H. (ed.) Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference, Santa Barbara, California, USA, August 23-27, 1998, Proceedings. Lecture Notes in Computer Science, vol. 1462, pp. 26–45. Springer (1998)

- [4] Bellare, M., Namprempre, C.: Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In: Okamoto, T. (ed.) *Advances in Cryptology - ASIACRYPT 2000, 6th International Conference on the Theory and Application of Cryptology and Information Security*, Kyoto, Japan, December 3-7, 2000, Proceedings. *Lecture Notes in Computer Science*, vol. 1976, pp. 531–545. Springer (2000)
- [5] Bellare, M., Namprempre, C.: Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *J. Cryptol.* **21**(4), 469–491 (2008)
- [6] Bellare, M., Sahai, A.: Non-malleable encryption: Equivalence between two notions, and an indistinguishability-based characterization. In: Wiener, M.J. (ed.) *CRYPTO*. vol. 1666, pp. 519–536. Springer (1999)
- [7] Boldyreva, A., Fischlin, M.: Analysis of random oracle instantiation scenarios for OAEP and other practical schemes. In: Shoup, V. (ed.) *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference*, Santa Barbara, California, USA, August 14-18, 2005, Proceedings. *Lecture Notes in Computer Science*, vol. 3621, pp. 412–429. Springer (2005)
- [8] Cramer, R., Shoup, V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk, H. (ed.) *Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference*, Santa Barbara, California, USA, August 23-27, 1998, Proceedings. *Lecture Notes in Computer Science*, vol. 1462, pp. 13–25. Springer (1998)
- [9] Dolev, D., Dwork, C., Naor, M.: Non-malleable cryptography (extended abstract). In: Koutsougeras, C., Vitter, J.S. (eds.) *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, May 5-8, 1991, New Orleans, Louisiana, USA. pp. 542–552. ACM (1991)
- [10] Duman, J., Hövelmanns, K., Kiltz, E., Lyubashevsky, V., Seiler, G.: Faster lattice-based kems via a generic fujisaki-okamoto transform using prefix hashing. In: Kim, Y., Kim, J., Vigna, G., Shi, E. (eds.) *CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security*, Virtual Event, Republic of Korea, November 15 - 19, 2021. pp. 2722–2737. ACM (2021)
- [11] Faust, S., Kohlweiss, M., Marson, G.A., Venturi, D.: On the non-malleability of the fiat-shamir transform. In: Galbraith, S.D., Nandi, M. (eds.) *Progress in Cryptology - INDOCRYPT 2012, 13th International Conference on Cryptology in India*, Kolkata, India, December 9-12, 2012. Proceedings. *Lecture Notes in Computer Science*, vol. 7668, pp. 60–79. Springer (2012)
- [12] Fischlin, M.: Completely non-malleable schemes. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) *Automata, Languages and Programming, 32nd International Colloquium, ICALP 2005*, Lisbon, Portugal, July 11-15, 2005, Proceedings. *Lecture Notes in Computer Science*, vol. 3580, pp. 779–790. Springer (2005)
- [13] Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener, M.J. (ed.) *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference*, Santa Barbara, California, USA, August 15-19, 1999, Proceedings. *Lecture Notes in Computer Science*, vol. 1666, pp. 537–554. Springer (1999)
- [14] Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the fujisaki-okamoto transformation. In: Kalai, Y., Reyzin, L. (eds.) *Theory of Cryptography - 15th International Conference, TCC 2017*, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part I. *Lecture Notes in Computer Science*, vol. 10677, pp. 341–371. Springer (2017)
- [15] Libert, B., Yung, M.: Efficient completely non-malleable public key encryption. In: Abramsky, S., Gavioille, C., Kirchner, C., auf der Heide, F.M., Spirakis, P.G. (eds.) *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010*, Bordeaux, France, July 6-10, 2010, Proceedings, Part I. *Lecture Notes in Computer Science*, vol. 6198, pp. 127–139. Springer (2010)

- [16] Nagao, W., Manabe, Y., Okamoto, T.: On the equivalence of several security notions of key encapsulation mechanism. IACR Cryptol. ePrint Arch. p. 268 (2006), <http://eprint.iacr.org/2006/268>
- [17] NIST: Nist announces first four quantum-resistant cryptographic algorithms (2022), <https://www.nist.gov/news-events/news/2022/07/nist-announces-first-four-quantum-resistant-cryptographic-algorithms>
- [18] Pass, R., Shelat, A., Vaikuntanathan, V.: Relations among notions of non-malleability for encryption. In: Kurosawa, K. (ed.) ASIACRYPT. vol. 4833, pp. 519–535 (2007)
- [19] Rackoff, C., Simon, D.R.: Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In: Feigenbaum, J. (ed.) Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings. Lecture Notes in Computer Science, vol. 576, pp. 433–444. Springer (1991)
- [20] Sepahi, R., Steinfeld, R., Pieprzyk, J.: Lattice-based completely non-malleable PKE in the standard model (poster). In: Parampalli, U., Hawkes, P. (eds.) Information Security and Privacy - 16th Australasian Conference, ACISP 2011, Melbourne, Australia, July 11-13, 2011. Proceedings. Lecture Notes in Computer Science, vol. 6812, pp. 407–411. Springer (2011)
- [21] Sepahi, R., Steinfeld, R., Pieprzyk, J.: Lattice-based completely non-malleable public-key encryption in the standard model. Des. Codes Cryptogr. **71**(2), 293–313 (2014)
- [22] Ventre, C., Visconti, I.: Completely non-malleable encryption revisited. In: Cramer, R. (ed.) Public Key Cryptography - PKC 2008, 11th International Workshop on Practice and Theory in Public-Key Cryptography, Barcelona, Spain, March 9-12, 2008. Proceedings. Lecture Notes in Computer Science, vol. 4939, pp. 65–84. Springer (2008)