Relations between Constrained and Bounded Chosen Ciphertext Security for Key Encapsulation Mechanisms

Takahiro Matsuda^{1,*}, Goichiro Hanaoka¹, and Kanta Matsuura²

¹ National Institute of Advanced Industrial Science and Technology, Japan
`{t-matsuda,hanaoka-goichiro}@aist.go.jp`

² The University of Tokyo, Japan
`kanta@iis.u-tokyo.ac.jp`

Abstract. In CRYPTO 2007, Hofheinz and Kiltz formalized a security notion for key encapsulation mechanisms (KEMs), called *constrained chosen ciphertext* (CCCA) security, which is strictly weaker than ordinary chosen ciphertext (CCA) security, and showed a new composition paradigm for CCA secure hybrid encryption. Thus, CCCA security of a KEM turned out to be quite useful. However, since the notion is relatively new and its definition is slightly complicated, relations among CCCA security and other security notions have not been clarified well. In this paper, in order to better understand CCCA security and the construction of CCCA secure KEMs, we study relations between CCCA and *bounded CCA* security, where the latter notion considers security against adversaries that make a-priori bounded number of decapsulation queries, and is also strictly weaker than CCA security. Specifically, we show that in most cases there are separations between these notions, while there is some unexpected implication from (a slightly stronger version of) CCCA security to a weak form of 1-bounded CCA security. We also revisit the construction of a KEM from a hash proof system (HPS) with computational security properties, and show that the HPS-based KEM, which was previously shown CCCA secure, is actually 1-bounded CCA secure as well. This result, together with the above general implication, suggests that 1-bounded CCA security can be essentially seen as a “necessary” condition for a CCCA secure KEM.

Keywords: key encapsulation mechanism, constrained CCA security, bounded CCA security, hash proof system.

1 Introduction

Background and Motivation. Studies on constructing and understanding practical public key encryption (PKE) schemes secure against chosen ciphertext attacks (CCA) [26,11] are important research topics in the area of cryptography.

* Takahiro Matsuda is supported by JSPS Research Fellowships for Young Scientists.

Among several approaches towards CCA secure PKE schemes, one of the promising approaches is to construct a PKE scheme via the hybrid encryption methodologies using a key encapsulation mechanism (KEM) which encapsulates (i.e. encrypts) a random session-key, and a data encapsulation mechanism (DEM) which encrypts an actual message using the session-key. Cramer and Shoup [10] show that if we combine a CCA secure KEM and a CCA secure DEM, we obtain a hybrid PKE scheme which is CCA secure. Abe et al. [1] show yet another hybrid encryption paradigm from a Tag-KEM, which is an extension of a KEM, and a passively secure DEM.

In CRYPTO 2007, Hofheinz and Kiltz [17] formalized a security notion for KEMs called *constrained chosen ciphertext* (CCCA) security, which is strictly weaker than ordinary CCA security. Then, they show that a CCA secure PKE scheme can be constructed by combining a CCCA secure KEM and a DEM satisfying the security of (one-time) authenticated encryption [4]. Therefore, CCCA security turned out to be a quite useful security notion for constructing a CCA secure PKE scheme.

However, the notion of CCCA security is relatively new, and the definition of CCCA security is slightly technically complicated compared to other existing security notions for KEMs, such as (ordinary) CCA security. Therefore, the relations between CCCA security and other security notions have not been studied and clarified well. Especially, “how” CCCA security is weak, compared to ordinary CCA security, seems not to have been understood well previously. It is naturally expected that the better we understand CCCA security itself, the higher the possibility we will come up with practical CCCA secure KEMs becomes, which will also lead to practical CCA secure PKE schemes.

So far, there are several positive and negative results regarding how close CCCA security and CCA security for KEMs are: Baek et al. [2] show that a CCCA secure KEM can be generically converted into a CCA secure one by using a one-time secure message authentication code. Hanaoka and Kurosawa [13] show that in fact, a CCCA secure KEM can be turned into a CCA secure one even without using any other additional building block, by using a part of the session-key (for a DEM) to check the consistency of a ciphertext in the decapsulation process. These results make us believe that CCCA security and CCA security for KEMs are in fact very close.

On the other hand, Choi et al. [7] show that the well-known KEM by Kurosawa and Desmedt [21], which was shown to be CCCA secure under the decisional Diffie-Hellman (DDH) assumption in [17], is not OW-2-CCA secure. That is, the session-key hidden in a ciphertext of the Kurosawa-Desmedt KEM can be recovered if an adversary can submit two decapsulation queries of its choice. This result, in contrast to the above positive results [2,13], makes us think that CCCA security is far from CCA security.

These previous results may illustrate that it is difficult to grasp what is actually achieved by CCCA security and what is not. The motivation of this work is to clarify the relations between CCCA security and other security notions, so that it leads to better understanding of CCCA security itself and also leads to

insights for constructing practical CCCA secure KEMs in the future. For that purpose, we study relations between CCCA security and *bounded CCA security* [8], which only captures security against adversaries that make a-priori bounded number of decapsulation queries (denoted by “ q -CCA” for q queries), and is also strictly weaker than CCA security in a different sense from CCCA security.

It is known that we can construct a “ q -bounded CCA” secure KEM whose ciphertext consists of only one group element (and thus “optimal” ciphertext size as a KEM) under the DDH assumption [8], for any predetermined polynomial q . On the other hand, the best known CCCA secure KEMs under the DDH assumption (or weaker assumptions) [21,17,12] have at least two group elements in a ciphertext. If we can construct a CCCA secure KEM under the DDH assumption with just one group element ciphertext, it will lead to (by combining it with a DEM satisfying the security of authenticated encryption) the best DDH-based PKE scheme in terms of the ciphertext overhead¹, i.e. one group element plus the ciphertext overhead caused by the DEM, which can be as small as k -bit for k -bit security. We believe that studying relations between CCCA security and bounded CCA security will also lead to important insights for the possibility of such “space-efficient” CCA secure PKE schemes (under DDH and weaker assumptions). Especially, understanding “how hard” it is to construct a CCCA secure KEM compared to a KEM with bounded CCA security will benefit the future designers of CCCA secure KEMs.

Our Contribution. Firstly, in Section 3 we investigate relations between CCCA security and bounded CCA security, i.e. implications/separations between these two security notions. One might expect that there is always a separation of CCCA security from bounded CCA security, and vice versa. As expected, we show that in most cases we have separations in both directions, and thus our contribution regarding this result is to give formal proofs, together with some basic ideas, for the separations. In particular, we show that IND-CCCA security does not imply OW-2-CCA or IND-1-CCA security (here, OW and IND stand for “one-wayness” and “indistinguishability”, respectively, and the formal definitions for security notions are given in Section 2). Perhaps somewhat surprisingly, however, it turns out that there is an implication from CCCA security to the weakest form of bounded CCA security, namely, OW-1-CCA, *if we slightly change the definition of a valid CCCA adversary*. The change we make to show the implication is regarding the definition of *uncertainty* that plays an important role in defining CCCA security, but is quite subtle. The proof for this result involves some unusual treatment (at least in the context of security proofs of CCCA/bounded CCA security) of an adversary, and might be of independent interest. For more details, see Section 3.2. We also show the separation of the opposite direction (bounded CCA security does not imply CCCA security) in terms of the number of queries allowed for an adversary. Specifically, we show that for any polynomial $q \geq 0$, IND- q -CCA security does not imply OW- $(q+1)$ -CCCA security.

¹ Ciphertext overhead is the difference between the size of a ciphertext and the size of its plaintext.

Then, in Section 4 we revisit the construction of KEMs from a *hash proof system* (HPS) [9,17], and show that the HPS-based KEM, which was shown to be IND-CCCA secure [17] (under some computational security requirements), actually satisfies IND-1-CCA security under the same assumptions used to show its IND-CCCA security. This result should be contrasted with the above mentioned separation of IND-1-CCA from IND-CCCA. Given the hybrid encryption paradigm by Cramer and Shoup [10], the result here implies that if we combine a HPS-based KEM (e.g. the Kurosawa-Desmedt KEM [21]) with a CCA secure redundancy-free DEM (e.g. a strong pseudorandom permutation [25]), the resulting PKE scheme still provides IND-1-CCA security. (As mentioned above, OW-2-CCA attack on this KEM is possible, and thus this is the optimal security result for the Kurosawa-Desmedt KEM, in terms of bounded CCA security.) Given the fact that using computational HPS is one of the major methodologies for constructing a practical CCCA secure KEM, we see that IND-1-CCA security can essentially be viewed as a “necessary” condition for a CCCA secure KEM. This result on the HPS-based KEM, together with the above general implication to OW-1-CCA security, also suggests that constructing a CCCA secure KEM is harder than constructing a 1-bounded CCA secure one. To the best of our knowledge, such insights have not been known previously.

Although it might be hard to imagine that 1-bounded CCA security (i.e. OW-1-CCA security and IND-1-CCA security) plays a practical role in real world applications in which KEMs (and PKE schemes) are used,² we stress that our aim in this paper is not to emphasize the importance of such security in practice, but rather to give better understanding of CCCA security itself, and we believe that our results give insights for constructing CCCA secure KEMs, and are useful for the future users/designers of CCCA secure KEMs.

Due to space limitation, the full proofs of the theorems in this paper will be given in the full version. We instead give proof sketches for each theorem.

Related Work. After Hofheinz and Kiltz [17] defined CCCA security, several practical CCCA secure KEMs have been proposed [6,12,20,13,14]. Hiwatari et al. [16] extended the CCCA secure KEM by Hanaoka and Kurosawa [12] to a CCCA secure multi-recipient KEM. Sakai et al. [27] used a OW-CCCA secure KEM which has reproducibility to construct a CCA secure KEM whose ciphertext length is shorter than that of the building block KEM, using a random oracle.

Bellare et al. [3] formalized the security notions for PKE schemes in a systematic way and showed the relations among security notions. For KEMs, Nagao et al. [23] and Herranz et al. [15] showed the relations among security notions. Moreover, Herranz et al. investigated the security notions achieved by hybrid encryption from a KEM and a DEM with several different levels of security.

² Very recently, Hohenberger et al. [18] used an IND-1-CCA secure PKE scheme as one of building blocks to construct a (fully) CCA secure PKE scheme. Although their construction still does not yield a practical scheme (at least compared to the concrete schemes, e.g. [10,21,17,6,12]), it would be interesting to seek for another application of 1-(or more-)bounded CCA secure schemes.

Bellare and Sahai [5] and later Pass, Shelat, and Vaikuntanathan [24] investigated the relations among several types of non-malleability [11]. Cramer et al. [8] introduced bounded CCA security, and show that non-malleability is separated from bounded CCA security. Matsuda and Matsuura [22] considered parallel decryption queries (which was originally introduced by Bellare and Sahai [5] in the context of non-malleability of PKE schemes) in bounded CCA security of PKE schemes and KEMs, and show several general implication/separation results. We note that the relations among security notions we show in this paper are not covered by these previous works.

2 Preliminaries

In this section, we review the basic notation and the definitions for a KEM.

Basic Notation. $\mathbb{N}$ denotes the set of all natural numbers, and if $n \in \mathbb{N}$ then $[n] = \{1, \dots, n\}$. “ $x \leftarrow y$ ” denotes that x is chosen uniformly at random from y if y is a finite set, x is output from y if y is a function or an algorithm, or y is assigned to x otherwise. If x and y are strings, then “ $|x|$ ” denotes the bit-length of x , “ $\text{msb}(x)$ ” denotes the most significant bit of x , and “ $x||y$ ” denotes a concatenation x and y . “PPTA” denotes a *probabilistic polynomial time algorithm*. If $\mathcal{A}$ is a probabilistic algorithm then $y \leftarrow \mathcal{A}(x; r)$ denotes that $\mathcal{A}$ computes y as output by taking x as input and using r as randomness. $\mathcal{A}^{\mathcal{O}}$ denotes an algorithm $\mathcal{A}$ with oracle access to $\mathcal{O}$. A function $f(k) : \mathbb{N} \rightarrow [0, 1]$ is said to be *negligible* if for all positive polynomials $p(k)$ and all sufficiently large $k \in \mathbb{N}$, we have $f(k) < 1/p(k)$.

Syntax of KEMs. A key encapsulation mechanism (KEM) Γ consists of the following three PPTAs (KG, Enc, Dec):

- KG:** The key generation algorithm that takes 1^k (security parameter k) as input, and outputs a public/secret key pair (pk, sk) .
- Enc:** The encapsulation algorithm that takes pk as input, and outputs a ciphertext c and a session-key $K \in \mathcal{K}$ (where $\mathcal{K}$ is the session-key space specified by pk).
- Dec:** The (deterministic) decapsulation algorithm that takes sk and c as input, and outputs a session-key K which could be a special symbol $\perp$ meaning “ c is an invalid ciphertext”.

We require $\text{Dec}(sk, c) = K$ for all (pk, sk) output by KG and all (c, K) output by $\text{Enc}(pk)$.

Security Notions for KEMs. Typically, security notions for KEMs are expressed by the combination of a security goal (GOAL) and an adversary’s attack type (ATK). In this paper, we will treat *indistinguishability* (IND) and *one-wayness* (OW) as security goals, and *chosen plaintext attacks* (CPA), *q-bounded chosen ciphertext attacks* (*q*-CCA) [8], *constrained CCA* (CCCA) [17], and its *q*-bounded analogue, namely, *q-bounded CCCA* (*q*-CCCA) as an adversary’s attack types, where $q \geq 0$ is an integer.

$\text{Expt}_{\Gamma, \mathcal{A}}^{\text{IND-ATK}}(k) :$ $(pk, sk) \leftarrow \text{KG}(1^k); b \leftarrow \{0, 1\};$ $(c^*, K_1^*) \leftarrow \text{Enc}(pk); K_0^* \leftarrow \mathcal{K};$ $b' \leftarrow \mathcal{A}^{\mathcal{O}}(pk, c^*, K_b^*);$ If $b' = b$ then return 1 else return 0	$\text{Expt}_{\Gamma, \mathcal{A}}^{\text{OW-ATK}}(k) :$ $(pk, sk) \leftarrow \text{KG}(1^k); (c^*, K^*) \leftarrow \text{Enc}(pk);$ $K' \leftarrow \mathcal{A}^{\mathcal{O}}(pk, c^*);$ If $K' = K^*$ then return 1 else return 0
---	---

Fig. 1. The security experiment for indistinguishability (IND-ATK experiment) (left) and that for one-wayness (OW-ATK experiment) (right)

For a KEM $\Gamma = (\text{KG}, \text{Enc}, \text{Dec})$, we define the experiment $\text{Expt}_{\Gamma, \mathcal{A}}^{\text{IND-ATK}}(k)$ in which an adversary $\mathcal{A}$ attacks indistinguishability of Γ under the attack type ATK, and the experiment $\text{Expt}_{\Gamma, \mathcal{A}}^{\text{OW-ATK}}(k)$ in which $\mathcal{A}$ attacks one-wayness of Γ under ATK, as in Fig. 1.

In the experiments, how the oracle $\mathcal{O}$ is defined and how it is available for $\mathcal{A}$ is determined depending on ATK in the following ways:

- If $\text{ATK} = \text{CPA}$, the oracle is unavailable and thus $\mathcal{A}$ cannot make any query.
- If $\text{ATK} = q\text{-CCA}$, the oracle is the decapsulation oracle $\mathcal{O}(\cdot) = \text{Dec}(sk, \cdot)$, and $\mathcal{A}$ can submit at most q queries. Furthermore, $\mathcal{A}$ is not allowed to submit the challenge ciphertext c^* to $\mathcal{O}$.
- If $\text{ATK} \in \{\text{CCCA}, q\text{-CCCA}\}$, the oracle is the *constrained decapsulation (CDEC) oracle* $\mathcal{O}_{\text{cdec}}(\cdot, \cdot)$, which takes a predicate $\text{pred} : \mathcal{K} \rightarrow \{0, 1\}$ and a ciphertext c as input, and returns a response as follows:

$$\mathcal{O}_{\text{cdec}}(\text{pred}(\cdot), c) = \begin{cases} K & \text{If } \text{Dec}(sk, c) = K \neq \perp \wedge \text{pred}(K) = 1 \\ \perp & \text{Otherwise} \end{cases}$$

Moreover, $\mathcal{A}$ is not allowed to submit a query containing c^* to $\mathcal{O}_{\text{cdec}}$. Additionally, if $\text{ATK} = q\text{-CCCA}$, $\mathcal{A}$ can submit at most q queries (as in $q\text{-CCA}$).

For a KEM Γ and $\text{GOAL} \in \{\text{IND}, \text{OW}\}$, let $\mathcal{A}$ be an adversary that runs in $\text{Expt}_{\Gamma, \mathcal{A}}^{\text{GOAL-CCCA}}(k)$ and makes in total Q queries, and let $(\text{pred}_i(\cdot), c_i)$ be $\mathcal{A}$'s i -th CDEC query. “*The running time of $\mathcal{A}$ in the GOAL-CCCA experiment*” is defined as the sum of $\mathcal{A}$'s running time and the total of maximum running time for evaluating each pred_i submitted by $\mathcal{A}$. “*The running time of the GOAL-CCCA experiment*” is defined as the total running time of the whole experiment $\text{Expt}_{\Gamma, \mathcal{A}}^{\text{GOAL-CCCA}}(k)$ minus “the running time of $\mathcal{A}$ in the GOAL-CCCA experiment”. For a CCCA adversary $\mathcal{A}$ and an experiment $\mathcal{E}$ (not necessarily $\text{Expt}_{\Gamma, \mathcal{A}}^{\text{GOAL-CCCA}}(k)$) that $\mathcal{A}$ runs in, we define the parameter called (plaintext) *uncertainty* $\text{uncert}_{\mathcal{A}, \mathcal{E}}(k)$ by:

$$\text{uncert}_{\mathcal{A}, \mathcal{E}}(k) = \frac{1}{Q} \sum_{i \in [Q]} \Pr[\mathcal{E}; K \leftarrow \mathcal{K} : \text{pred}_i(K) = 1]$$

Finally, we say that an adversary $\mathcal{A}$ is a *valid GOAL-CCCA adversary* if (1) “the running time of $\mathcal{A}$ in the GOAL-CCCA experiment” is polynomial in k , and (2) $\text{uncert}_{\mathcal{A}, \mathcal{E}}(k)$ is negligible for all experiments $\mathcal{E}$ whose running time is at most

“the running time of GOAL-CCCA experiment” that $\mathcal{A}$ runs in. We define the notion of a “valid GOAL- q -CCCA adversary” in exactly the same way as above.

For a KEM Γ , an adversary $\mathcal{A}$, and $\text{ATK} \in \{\text{CPA}, q\text{-CCA}, \text{CCCA}, q\text{-CCCA}\}$, we define IND-ATK advantage $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{IND-ATK}}(k)$ and OW-ATK advantage $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{OW-ATK}}(k)$ by:

$$\begin{aligned}\text{Adv}_{\Gamma, \mathcal{A}}^{\text{IND-ATK}}(k) &= |\Pr[\text{Expt}_{\Gamma, \mathcal{A}}^{\text{IND-ATK}}(k) = 1] - \frac{1}{2}| \\ \text{Adv}_{\Gamma, \mathcal{A}}^{\text{OW-ATK}}(k) &= \Pr[\text{Expt}_{\Gamma, \mathcal{A}}^{\text{OW-ATK}}(k) = 1]\end{aligned}$$

Definition 1. Let $\text{GOAL} \in \{\text{IND}, \text{OW}\}$ and $q \in \mathbb{N}$. We say that a KEM Γ is GOAL-CPA (resp. GOAL- q -CCA) secure if $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{GOAL-CPA}}(k)$ (resp. $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{GOAL-}q\text{-CCA}}(k)$) is negligible for any PPTA $\mathcal{A}$. We say that a KEM Γ is GOAL-CCCA (resp. GOAL- q -CCCA) secure if $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{GOAL-CCCA}}(k)$ (resp. $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{GOAL-}q\text{-CCCA}}(k)$) is negligible for any valid GOAL-CCCA (resp. GOAL- q -CCCA) adversary $\mathcal{A}$.

3 Relations between Constrained and Bounded Chosen Ciphertext Security

In this section, we investigate relations between constrained and bounded CCA security. One might expect that there is always a separation of CCCA security from bounded CCA security, and vice versa. It is actually the case, and we formally show that for most cases we have separations in both directions. Perhaps somewhat surprisingly, however, it turns out that there is an implication from IND-CCCA security to the weakest form of bounded CCA security, namely, OW-1-CCA, if we slightly change the definition of a valid CCCA adversary.

The rest of this section is organized as follows: In Section 3.1, we show the separations between CCCA and bounded CCA security. Then, in Section 3.2 we introduce slightly stronger CCCA security and its implication to OW-1-CCA security.

3.1 Separations

Basic Ideas for Separations. Notice that a CDEC query by a valid CCCA adversary $\mathcal{A}$ is answered with a value that is not $\perp$ only when $\mathcal{A}$ already has some “non-trivial” knowledge about the decapsulation result, where the non-triviality is captured by the condition that $\mathcal{A}$ has to control the uncertainty negligible. We note that CDEC queries made by a valid CCCA adversary $\mathcal{A}$ cannot (except with negligible probability) reveal information on the decapsulation result that is hard to guess and is independent from $\mathcal{A}$ ’s view, because otherwise $\mathcal{A}$ ’s uncertainty cannot be negligible. We use this idea for showing the separations of bounded CCA security from CCCA security.

On the other hand, CDEC queries by a valid adversary $\mathcal{A}$ can reveal (while controlling $\mathcal{A}$ ’s uncertainty negligible) information that is dependent on some part of a public key, even if the decapsulation result itself is hard to guess, as

$\text{KG}_{\text{sep1}}(1^k) :$ $(pk, sk) \leftarrow \text{KG}(1^k)$ $R \leftarrow \{0, 1\}^k$ $SK \leftarrow (sk, R)$ Return (pk, SK) .	$\text{Enc}_{\text{sep1}}(pk) :$ $(c, K) \leftarrow \text{Enc}(pk)$ $C \leftarrow (0 c)$ Return (C, K) .	$\text{KG}_{\text{sep2}}(1^k) :$ $(pk, sk) \leftarrow \text{KG}(1^k)$ $R \leftarrow \{0, 1\}^{k-1}$ $SK \leftarrow (sk, R)$ Return (pk, SK) .	$\text{Enc}_{\text{sep2}}(pk) :$ $(c, K) \leftarrow \text{Enc}(pk)$ $C \leftarrow (0 c)$ Return (C, K) .
$\text{Dec}_{\text{sep1}}(SK, C) :$ Parse SK as (sk, R) and C as (γc) . $K \leftarrow \text{Dec}(sk, c)$ If $\gamma = 0$ or $K = \perp$ then return K . Return $K \oplus R$.	$\text{Dec}_{\text{sep2}}(SK, C) :$ Parse SK as (sk, R) and C as (γc) . $K \leftarrow \text{Dec}(sk, c)$ If $\gamma = 0$ or $K = \perp$ then return K . Return $(\text{msb}(K) R)$.		
$\text{KG}_{\text{sep3}}(1^k) :$ $R_{\text{KG}} \leftarrow \{0, 1\}^k$ $(pk, sk) \leftarrow \text{KG}(1^k; R_{\text{KG}})$ $v_0 \leftarrow 1^k$ If $q \geq 1$ then $v_i \leftarrow \{0, 1\}^k$ for $i \in [q]$ $V_i \leftarrow f(v_i)$ for $i \in \{0, \dots, q\}$ $PK \leftarrow (pk, \{V_i\}_{i \in \{0, \dots, q\}})$ $SK \leftarrow (sk, \{v_i\}_{i \in \{0, \dots, q\}}, R_{\text{KG}})$ Return (PK, SK) .	$\text{Dec}_{\text{sep3}}(SK, C) :$ Parse SK as $(sk, \{v_i\}_{i \in \{0, \dots, q\}}, R_{\text{KG}})$. Parse C as (αc) s.t. $ \alpha = k$. Interpret α as an integer. If $\alpha = 0$ then return $K \leftarrow \text{Dec}(sk, c)$. If $\alpha \in [q]$ and $c = V_\alpha$ then return v_α . If $\alpha = q + 1$ then Parse c as $(u_0 u_1 \dots u_q)$ s.t. $ u_i = k$ for $i \in \{0, \dots, q\}$. (If parsing fails then return $\perp$.) If $f(u_i) = V_i$ for all $i \in \{0, \dots, q\}$ then return R_{KG} . End if Return $\perp$.		
$\text{Enc}_{\text{sep3}}(PK) :$ Parse PK as $(pk, \{V_i\}_{i \in \{0, \dots, q\}})$. $(c, K) \leftarrow \text{Enc}(pk)$ $C \leftarrow (0^k c)$ Return (C, K) .			

Fig. 2. The KEM Γ_{sep1} that separates OW-2-CCA from IND-CCCA (upper-left), the KEM Γ_{sep2} that separates IND-1-CCA from IND-CCCA (upper-right), and the KEM Γ_{sep3} that separates OW- $(q+1)$ -CCCA from IND- q -CCA (bottom). In Γ_{sep3} , f is a one-way function.

long as it is non-trivial. This idea is later used to separate CCCA security from bounded CCA security in terms of the number of queries.

For simplicity, in this subsection we assume that the session-key space of a KEM is $\{0, 1\}^k$ when the key generation algorithm is run with input 1^k .

IND-CCCA vs. OW-2-CCA. Choi et al. [7] showed that the KEM part of the Kurosawa-Desmedt PKE scheme [21], which was shown to be IND-CCCA secure under the DDH assumption in [17], is not OW-2-CCA secure. This result implies that if there is a group with prime order in which the DDH assumption holds, then there exists a KEM which is IND-CCCA secure but is not OW-2-CCA secure (and thus IND-CCCA security does not imply OW-2-CCA security, under the DDH assumption). We remove the DDH assumption from this statement, and show that in general IND-CCCA security does not imply OW-2-CCA security.

Theorem 1. *If there exists an IND-CCCA secure KEM, then there exists a KEM which is IND-CCCA secure but is not OW-2-CCA secure. Moreover, the OW-2-CCA attack for the latter KEM succeeds even if an adversary has to make two decapsulation queries parallelly (i.e. non-adaptively).*

Proof Sketch. Let $\Gamma = (\text{KG}, \text{Enc}, \text{Dec})$ be an IND-CCCA secure KEM. Using the KEM Γ , we construct another KEM $\Gamma_{\text{sep1}} = (\text{KG}_{\text{sep1}}, \text{Enc}_{\text{sep1}}, \text{Dec}_{\text{sep1}})$ for the separation as in Fig. 2 (upper-left).

The OW-2-CCA attack for Γ_{sep1} is easy: Consider the following OW-2-CCA adversary $\mathcal{A}$. Given $(pk, C^* = (0||c^*))$, $\mathcal{A}$ computes $(c', K') \leftarrow \text{Enc}(pk)$, and submits ciphertexts $C_1 = (1||c^*)$ and $C_2 = (1||c')$ parallelly (i.e. non-adaptively) as decapsulation queries. According to the definition of Dec_{sep1} , $\mathcal{A}$ receives $K_1 = K^* \oplus R$ and $K_2 = K' \oplus R$, respectively, from the decapsulation oracle. Then $\mathcal{A}$ calculates $R \leftarrow K_2 \oplus K'$ and $K^* \leftarrow K_1 \oplus R$, and terminates with output K^* . It is easy to see that $\mathcal{A}$'s OW-2-CCA advantage is 1.

In order to show that Γ_{sep1} is IND-CCCA secure based on IND-CCCA security of the building block KEM Γ , consider the following sequence of games:

Game 1. This is the original IND-CCCA experiment, i.e. $\text{Expt}_{\Gamma_{\text{sep1}}, \mathcal{A}}^{\text{IND-CCCA}}(k)$.

Game 2. Same as Game 1, except that any CDEC query containing a ciphertext of the form $C = (1||c)$ is answered with $\perp$.

Let $\mathcal{A}$ be any valid IND-CCCA adversary that makes in total Q CDEC queries. Then the difference in Game 1 and Game 2 can occur only when $\mathcal{A}$ submits a CDEC query (pred, C) satisfying $C = (1||c)$, $\text{Dec}_{\text{sep1}}(SK, C) = K \neq \perp$, and $\text{pred}(K) = 1$. (In Game 1, it is answered with K , while in Game 2 it is answered with $\perp$.) By definition of Dec_{sep1} , if $C = (1||c)$ and $\text{Dec}(sk, c) \neq \perp$, then $\text{Dec}_{\text{sep1}}(SK, C) = \text{Dec}(sk, c) \oplus R$. However, notice that the information on R is information-theoretically hidden from $\mathcal{A}$'s view in Game 2. Moreover, R is chosen uniformly from $\{0, 1\}^k$, and thus the decapsulation result $\text{Dec}(sk, c) \oplus R$ of the query of the above type is also uniformly random and independent of $\mathcal{A}$'s view in Game 2. Then, the probability that some of $\mathcal{A}$'s CDEC queries of the form $(\text{pred}, C = (1||c))$ satisfies $\text{Dec}_{\text{sep1}}(SK, C) = K' \neq \perp$ and $\text{pred}(K') = 1$ will be upperbounded by $Q \cdot \text{uncert}_{\mathcal{A}, \text{Game 2}}(k)$, which is negligible due to the fact that $\mathcal{A}$ is a valid IND-CCCA adversary. Moreover, Game 2 can be perfectly simulated by another valid IND-CCCA adversary for the building block KEM Γ , which means that $\mathcal{A}$'s advantage in Game 2 is negligible. In summary, $\mathcal{A}$'s IND-CCCA advantage is upperbounded to be negligible. $\square$

IND-CCCA vs. Non-malleability. In the above theorem, to break OW-2-CCA security of the KEM Γ_{sep1} , the two decapsulation queries can be made parallelly. Hence, due to the equivalence of non-malleability under chosen plaintext attack and indistinguishability under one parallel decapsulation query [5,23,15], and the transitivity of the implication of security notions, it follows that IND-CCCA security does not imply non-malleability (under chosen plaintext attack).

IND-CCCA vs. IND-1-CCA. We next show that if the security goal is IND, then even IND-1-CCA security is in general separated from IND-CCCA.

Theorem 2. *If there exists an IND-CCCA secure KEM, then there exists a KEM which is IND-CCCA secure but is not IND-1-CCA secure.*

Proof Sketch. Let $\Gamma = (\text{KG}, \text{Enc}, \text{Dec})$ be an IND-CCCA secure KEM. Using the KEM Γ , we construct another KEM $\Gamma_{\text{sep2}} = (\text{KG}_{\text{sep2}}, \text{Enc}_{\text{sep2}}, \text{Dec}_{\text{sep2}})$ for the separation as in Fig. 2 (upper-right).

The IND-1-CCA attack for Γ_{sep2} is quite easy to see. Consider the following IND-1-CCA adversary $\mathcal{A}$. Given $(pk, C^* = (0||c^*), K_b^*)$, $\mathcal{A}$ submits a decapsulation query $C = (1||c^*)$, and receives the result K , which must be of the form $K = (\text{msb}(K_1^*)||R)$ according to the definition of Dec_{sep2} . Then $\mathcal{A}$ checks if $\text{msb}(K_b^*) = \text{msb}(K)$, and outputs $b' = 1$ if this is the case, and outputs $b' = 0$ otherwise. A simple calculation shows that $\mathcal{A}$'s IND-1-CCA advantage is $1/4$.

The proof of IND-CCCA security of Γ_{sep2} based on IND-CCCA security of the building block KEM Γ proceeds almost in the same way as that of Γ_{sep1} , considering the two games Game 1 ($\text{Expt}_{\Gamma_{\text{sep2}}, \mathcal{A}}^{\text{IND-CCCA}}(k)$ itself) and Game 2 (in which every CDEC query containing a ciphertext of the form $C = (1||c)$ is rejected). Game 1 and Game 2 are identical unless a valid IND-CCCA adversary $\mathcal{A}$ makes a CDEC query (pred, C) satisfying $C = (1||c)$, $\text{Dec}_{\text{sep2}}(SK, C) = K \neq \perp$, and $\text{pred}(K) = 1$. The decapsulation result of such a ciphertext is of the form $(\text{msb}(\text{Dec}(sk, c))||R)$ where R is the value in the secret key (if c is not invalid). However, recall that this R is chosen uniformly at random, and is information-theoretically hidden from $\mathcal{A}$ and independent of $\mathcal{A}$'s view in Game 2. Therefore, the predicates contained in $\mathcal{A}$'s CDEC queries of the above type are almost never satisfied by the corresponding decapsulation results due to the condition that $\mathcal{A}$ has to control its uncertainty negligible, which implies that the difference between $\mathcal{A}$'s success probability (in guessing the challenge bit) in Game 1 and that in Game 2 is negligible. More specifically, in the full proof, we show that the difference in $\mathcal{A}$'s success probability in these games is upperbounded by $2Q \cdot \text{uncert}_{\mathcal{A}, \text{Game 2}}(k)$ where Q is the total number of $\mathcal{A}$'s CDEC queries (the reason why “2” appears is because the value R in Γ_{sep2} is not k -bit but $(k-1)$ -bit, and we lose the factor 2 when relating it with the uncertainty that considers whether the predicates are satisfied by a k -bit randomness). The fact that $\mathcal{A}$'s advantage in Game 2 is negligible follows from the IND-CCCA security of the building block KEM Γ , as in Γ_{sep1} . $\square$

IND- q -CCA vs. OW- $(q+1)$ -CCCA. The above separations show that CCCA security does not imply bounded CCA security in most cases. Here, we show the separation of the opposite direction: if there is no trivial implication in terms of the number of queries, CCCA security is separated from bounded CCA security.

Theorem 3. *For any polynomial $q \geq 0$, if there exists an IND- q -CCA secure KEM, then there exists a KEM which is IND- q -CCA secure but is not OW- $(q+1)$ -CCCA secure.*

Proof Sketch. Fix $q \geq 0$. Let $\Gamma = (\text{KG}, \text{Enc}, \text{Dec})$ be an IND- q -CCA secure KEM. Here, without loss of generality, we assume that the randomness space of KG is $\{0, 1\}^k$. Moreover, let $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be a one-way function (OWF), whose existence is guaranteed by the existence of Γ . Using the KEM Γ and the OWF f , we construct another KEM $\Gamma_{\text{sep3}} = (\text{KG}_{\text{sep3}}, \text{Enc}_{\text{sep3}}, \text{Dec}_{\text{sep3}})$ for the

separation as in Fig. 2 (bottom). In the following, whenever we treat an integer as a k -bit string, we use “hat” (e.g. $\widehat{1}$ is the k -bit representation of 1).

The OW- $(q+1)$ -CCA attack against the KEM Γ_{sep3} is as follows (if $q = 0$, then we skip this part and goes to the $(q+1)$ -th query below): Given $(PK, C^* = (0^k || c^*))$, for $i \in [q]$, a OW- $(q+1)$ -CCCA adversary $\mathcal{A}$ defines the predicate pred_i by “ $\text{pred}_i(K) = 1$ iff $f(K) = V_i$ ”, and submits the i -th CDEC query of the form $(\text{pred}_i, C_i = (\widehat{i} || V_i))$. Since $\text{Dec}_{\text{sep3}}(SK, C_i) = v_i$ by definition, $\mathcal{A}$ receives v_i from the oracle. After obtaining $v_1, \dots, v_q$, $\mathcal{A}$ defines the $(q+1)$ -th predicate pred_{q+1} by “ $\text{pred}_{q+1}(K) = 1$ iff $\text{KG}(1^k; K) = (pk, *)$ ”, sets $C_{q+1} \leftarrow (\widehat{q+1} || v_0 || v_1 || \dots || v_q)$, and submits $(\text{pred}_{q+1}, C_{q+1})$ to the oracle as the $(q+1)$ -th CDEC query. Since $\text{Dec}(SK, C_{q+1}) = R_{\text{KG}}$, $\mathcal{A}$ receives R_{KG} as a response. $\mathcal{A}$ can then compute sk from R_{KG} and decrypt c^* , and thus $\mathcal{A}$ ’s OW- $(q+1)$ -CCCA advantage is 1.

Here, we also have to show that the above $\mathcal{A}$ is a valid OW- $(q+1)$ -CCCA adversary. We have to be careful because we have to show that $\mathcal{A}$ ’s uncertainty is negligible for *any* experiment $\mathcal{E}$ that is as efficient as the original OW- $(q+1)$ -CCCA experiment.³ Fortunately, we can use the following statistical property that is satisfied by any OWF. (The proof is given in the full version.)

Lemma 1. *If f is a OWF, then $\Pr[x \leftarrow \{0, 1\}^k : f(x) = y]$ is negligible for any string $y \in \{0, 1\}^*$.*

This guarantees that, for $i \in [q]$, whatever value is assigned to V_i by an experiment $\mathcal{E}$, the probability that pred_i is satisfied by a random K is negligible. Furthermore, recall that the key generation algorithm of any secure (at least OW-CPA secure) KEM can be viewed as a OWF whose domain is the randomness space of KG and whose image is pk (sk is discarded). Then we can use Lemma 1 also for the $(q+1)$ -th CDEC predicate pred_{q+1} , and conclude that $\text{uncert}_{\mathcal{A}, \mathcal{E}}(k)$ is negligible for any experiment $\mathcal{E}$.

IND- q -CCA security of Γ_{sep3} is explained as follows. Let $\mathcal{A}$ be any IND- q -CCA adversary against Π_{sep3} . Recall that a decapsulation query of the form $C = (\widehat{q+1} || c)$ is answered with R_{KG} only when all preimages $v_0, v_1, \dots, v_q$ are known to $\mathcal{A}$. Since v_0 is the fixed value 1^k , $\mathcal{A}$ actually needs to find q preimages $v_1, \dots, v_q$. However, due to one-wayness of f , it is hard to find v_i without making a decapsulation query of the form $C = (\widehat{i} || V_i)$. But since $\mathcal{A}$ can make only q queries, if $\mathcal{A}$ makes q queries to obtain $(v_1, \dots, v_q)$, $\mathcal{A}$ can no longer use the decapsulation oracle. This means that unless $\mathcal{A}$ breaks the OWF, $\mathcal{A}$ cannot make a decapsulation query of the form $C = (\widehat{q+1} || c)$ that reveals R_{KG} . Then, in order to break IND- q -CCA security of Γ_{sep3} , $\mathcal{A}$ has to essentially break IND- q -CCA security of the building block KEM Γ , which is hard by assumption. $\square$

3.2 Slightly Stronger CCCA Security and Its Implication

In the previous subsection, we have seen that IND-CCCA security does not imply OW-2-CCA or IND-1-CCA security. Then, a natural question would be whether

³ For example, $\mathcal{A}$ ’s uncertainty has to be negligible in which PK and/or C^* are generated incorrectly (as long as the experiment is efficient).

IND-CCCA security implies (or does not imply) OW-1-CCA security, which is the weakest bounded CCA security for KEMs. Actually, we could not show implication/separation from IND-CCCA. Alternatively, however, we find that if we consider a slightly stronger definition for IND-CCCA, we actually have an implication. The modification we will make is in the definition of uncertainty, and is quite subtle. We explain this in this subsection.

Note that the IND-CCCA experiment is fixed if we fix the following: (a) the randomness for key generation $((pk, sk) \leftarrow \text{KG}(1^k))$, (b) the randomness for challenge ciphertext/session-keys $((c^*, K_1^*) \leftarrow \text{Enc}(pk) \text{ and } K_0^* \leftarrow \mathcal{K})$, (c) the challenge bit $(b \leftarrow \{0, 1\})$, and (d) the randomness for an adversary. We denote the process of randomly picking these randomness and fixing the IND-CCCA experiment by $\mathcal{E} \leftarrow \text{Expt}_{\Gamma, \mathcal{A}}^{\text{IND-CCCA}}(k)$. We introduce the following definition.

Definition 2. Let Γ be a KEM and $\mathcal{A}$ be an IND-CCCA adversary (against Γ) that makes Q CDEC queries. Let pred_i be the predicate contained in $\mathcal{A}$'s i -th CDEC query. We define the average uncertainty $\text{uncert}_{\mathcal{A}}^{\text{ave}}(k)$ of $\mathcal{A}$ by:

$$\begin{aligned} \text{uncert}_{\mathcal{A}}^{\text{ave}}(k) &= \mathbf{E}_{\mathcal{E} \leftarrow \text{Expt}_{\Gamma, \mathcal{A}}^{\text{IND-CCCA}}(k)} [\text{uncert}_{\mathcal{A}, \mathcal{E}}(k)] \\ &= \mathbf{E}_{\mathcal{E} \leftarrow \text{Expt}_{\Gamma, \mathcal{A}}^{\text{IND-CCCA}}(k)} \left[\frac{1}{Q} \sum_{i \in [Q]} \Pr[\mathcal{E}; K \leftarrow \mathcal{K} : \text{pred}_i(K) = 1] \right] \end{aligned}$$

Furthermore, we say that $\mathcal{A}$ is a valid IND-CCCA* adversary (against Γ) if (1) the running time of $\mathcal{A}$ in the IND-CCCA experiment is polynomial in k , and (2) $\text{uncert}_{\mathcal{A}}^{\text{ave}}(k)$ is negligible.

Using average uncertainty, we define IND-CCCA* security of a KEM as follows:

Definition 3. We say that a KEM Γ is IND-CCCA* secure if $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{IND-CCCA}}(k)$ is negligible for any valid IND-CCCA* adversary $\mathcal{A}$.

We define OW-CCCA*, IND- q -CCCA*, and OW- q -CCCA* security in exactly the same way as above.

Note that to define IND-CCCA* security, we have not changed anything about the definition of IND-CCCA advantage $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{IND-CCCA}}(k)$. The only difference between IND-CCCA* security defined here and the original IND-CCCA security in [17] is for which class of adversaries we require the advantage to be negligible. In order for a CCA adversary $\mathcal{A}$ to be valid as an IND-CCCA* adversary, $\mathcal{A}$ only needs to control his uncertainty in the original IND-CCCA experiment to be negligible on an average, and thus for example, its uncertainty can be 1 accidentally (as long as it is negligible on an average). On the other hand, the original IND-CCCA security definition requires that $\mathcal{A}$'s uncertainty to be negligible for any experiment whose running time is at most that of the original IND-CCCA experiment. Therefore, if $\mathcal{A}$ is a valid IND-CCCA adversary, then it is a valid IND-CCCA* adversary as well. Since IND-CCCA security requires the IND-CCCA advantage to be negligible for adversaries of a smaller class, IND-CCCA* security implies IND-CCCA security.

Although the difference between IND-CCCA* and IND-CCCA security seems quite subtle and small, so far we are not sure if the latter implies (or is separated from) the former, and we would like to leave it as an open problem.

Now, we show the implication that bridges CCCA and bounded CCA security.

Theorem 4. *If a KEM is IND-1-CCCA* secure, then it is OW-1-CCA secure.*

Proof Sketch. Without loss of generality, a OW-1-CCA adversary $\mathcal{A}$ can be divided into two stages $(\mathcal{A}_1, \mathcal{A}_2)$ so that the OW-1-CCA experiment is rewritten as:

$(pk, sk) \leftarrow \text{KG}(1^k); (c^*, K^*) \leftarrow \text{Enc}(pk); (\hat{c}, \text{st}) \leftarrow \mathcal{A}_1(pk, c^*); \hat{K} \leftarrow \text{Dec}(sk, \hat{c});$
 $K' \leftarrow \mathcal{A}_2(\hat{K}, \text{st}); \text{ If } K' = K^* \text{ then return 1 else return 0}$

where $\hat{c}$ represents $\mathcal{A}$'s decapsulation query (which can be made only once). Moreover, we can assume that $\mathcal{A}_2$ is deterministic because in case $\mathcal{A}_2$ needs randomness, it can be chosen by $\mathcal{A}_1$ and passed via st . Now, using a OW-1-CCA adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, we construct an IND-1-CCCA adversary $\mathcal{B}$ as follows:

On input (pk, c^*, K_b^*) (where b is $\mathcal{B}$'s challenge bit), $\mathcal{B}$ runs $(\hat{c}, \text{st}) \leftarrow \mathcal{A}_1(pk, c^*)$. Then $\mathcal{B}$ defines pred by “ $\text{pred}(K) = 1$ iff $\mathcal{A}_2(K, \text{st}) = K_b^*$ ” and submits a CDEC query $(\text{pred}, \hat{c})$. If the answer from $\mathcal{O}_{cdec}$ is not $\perp$, $\mathcal{B}$ outputs 1. Otherwise $\mathcal{B}$ checks if $\mathcal{A}_2(\perp, \text{st}) = K_b^*$, and returns 1 if the check holds or returns 0 otherwise.

Assume that $\mathcal{A}$ breaks OW-1-CCA security with non-negligible advantage. Since $\mathcal{A}_2$, which is given a correct decapsulation $\text{Dec}(sk, \hat{c})$ during the evaluation of pred in $\mathcal{O}_{cdec}$, outputs $K_1^* = \text{Dec}(sk, c^*)$ with non-negligible probability, $\mathcal{B}$ outputs 1 with non-negligible probability when $b = 1$. (The check “ $\mathcal{A}_2(\perp, \text{st}) = K_b^*$ ” performed by $\mathcal{B}$ covers the case in which $\text{Dec}(sk, \hat{c}) = \perp$.) On the other hand, K_0^* is information-theoretically hidden from $\mathcal{A}$'s view, and $\mathcal{A}_2$ can output it only with negligible probability. Thus, when $b = 0$, pred is almost never satisfied and $\mathcal{B}$ outputs 1 only with negligible probability. Therefore, there is a non-negligible difference in the probabilities that $\mathcal{B}$ outputs 1 between the cases $b = 1$ and $b = 0$, namely, $\mathcal{B}$ has non-negligible IND-1-CCCA advantage. The idea of using an adversary in the predicate in a CDEC query might be of independent interest.

Note that $\mathcal{B}$'s uncertainty depends on $\mathcal{A}$, and we could not rule out the possibility that there is an experiment (which is as efficient as the IND-1-CCCA experiment) such that when $\mathcal{B}$ (which internally runs $\mathcal{A}$) is run, $\mathcal{B}$'s uncertainty is non-negligible. However, it is possible, using IND-1-CCCA* security of the KEM itself, to show that $\mathcal{B}$'s average uncertainty is negligible, and thus $\mathcal{B}$ is shown to be a valid IND-1-CCCA* adversary. We can show this roughly because the definition of average uncertainty considers the distribution of the public key and the challenge ciphertext/session-key pair (rather than fixed values for these), which makes it possible to use security of the KEM itself. Specifically, in the full proof we show that if the KEM is OW-CPA secure (which is trivially satisfied by the IND-1-CCCA* security of the KEM), then $\mathcal{B}$ is a valid IND-1-CCCA* adversary. $\square$

4 KEMs from Computational Hash Proof Systems, Revisited

In this section, we revisit the construction of KEMs from a HPS [17] for which we only require computational security properties, as opposed to the information-theoretic ones in [9,21]. More concretely, we show that if a HPS satisfies the computational security requirements defined in [17], then the KEM constructed based on the HPS satisfies not only IND-CCCA security but also IND-1-CCA security. (In particular, our result implies that the Kurosawa-Desmedt KEM [21] is IND-1-CCA secure under the DDH assumption.) This result should be contrasted with the separation of IND-1-CCA from IND-CCCA security in Section 3.2.

Below, we review the definitions of computational HPS in Section 4.1, and we show that the HPS-based KEM satisfies IND-1-CCA security in Section 4.2.

4.1 Definitions for Computational HPS

Here, we review the definition of hash proof systems as defined by Cramer and Shoup [9,21,17]. (We mainly borrow the notations from [17], which we customize slightly for our purpose.)

Let $\mathcal{C}$, $\mathcal{K}$, $\mathcal{S}$, and $\mathcal{P}$ be sets, and $\mathcal{V}$ be the set of “languages” satisfying $\mathcal{V} \subset \mathcal{C}$. Let $D_{sk} : \mathcal{C} \rightarrow \mathcal{K}$ be a hash function indexed by $sk \in \mathcal{S}$. Informally speaking, a HPS is a special type of a designated-verifier proof system for a “subset membership problem” (i.e. whether a “statement” $c \in \mathcal{C}$ satisfies $c \in \mathcal{V}$). A hash function D_{sk} is said to be *projective* if there exists an efficiently computable projection $\mu : \mathcal{S} \rightarrow \mathcal{P}$ such that $pk = \mu(sk) \in \mathcal{P}$ defines the action of D_{sk} over the subset $\mathcal{V}$. That is, for every $c \in \mathcal{V}$, the value $K = D_{sk}(c)$ is uniquely determined by $\mu(sk)$ and c . In the context of the HPS-based KEM that will be explained later, we will identify $\mathcal{C}$ as the ciphertext space, $\mathcal{V}$ as the set of all valid ciphertexts, $\mathcal{S}$ as the secret key space, $\mathcal{P}$ as the public key space, $\mathcal{K}$ as the session-key space, $\mu(\cdot)$ as the key generation algorithm, and $D_{sk}(\cdot)$ as the decapsulation algorithm. Taking this into account, hereafter we call an element $c \in \mathcal{C}$ *valid* if $c \in \mathcal{V}$ and *invalid* if $c \in \mathcal{C} \setminus \mathcal{V}$. As usual, we require: (1) $\mathcal{C}$ is efficiently recognizable, (2) a valid element $c \in \mathcal{V}$ can be efficiently sampled together with a witness w about the fact that $c \in \mathcal{V}$, and (3) we can sample elements from $\mathcal{C} \setminus \mathcal{V}$, $\mathcal{S}$, and $\mathcal{K}$ efficiently and (statistically close to) uniformly.

The above are the description of the parameters for a HPS. For simplicity, we assume that the definitions of the sets and the functions we described above are generated and determined by a probabilistic algorithm HGen. Formally, a HPS Π consists of the following three PPTAs (HGen, Pub, Priv):

HGen: The parameter generation algorithm for HPS which takes 1^k as input, and outputs parameters $\text{pub} = (\mathcal{C}, \mathcal{V}, \mathcal{K}, \mathcal{S}, \mathcal{P}, D_{(\cdot)} : \mathcal{C} \rightarrow \mathcal{K}, \mu : \mathcal{S} \rightarrow \mathcal{P})$.

For notational convenience, we assume that pub is provided as input to the following algorithms Pub and Priv, and do not write it explicitly.

Pub: The (deterministic) public evaluation algorithm which takes $pk = \mu(sk) \in \mathcal{P}$, a valid element/witness pair (c, w) (where w is about the fact that $c \in \mathcal{V}$) as input, and outputs a hash value $K = D_{sk}(c)$.

$\text{Expt}_{\Pi, \mathcal{A}}^{\text{CU}_2}(k):$ $\text{pub} \leftarrow \text{HGen}(1^k);$ $sk \leftarrow \mathcal{S}; pk \leftarrow \mu(sk);$ $c^* \leftarrow \mathcal{C} \setminus \mathcal{V}; K^* \leftarrow \mathcal{D}_{sk}(c^*);$ $(c', \text{st}) \leftarrow \mathcal{A}_1^{\mathcal{O}}(\text{pub}, pk, c^*, K^*);$ $K'_1 \leftarrow \mathcal{D}_{sk}(c'); K'_0 \leftarrow \mathcal{K}; b \leftarrow \{0, 1\};$ $b' \leftarrow \mathcal{A}_2(K'_b, \text{st});$ If $b' = b$ then return 1 else return 0	$\text{Expt}_{\Pi, \mathcal{A}}^{\text{CU}_1}(k):$ $\text{pub} \leftarrow \text{HGen}(1^k);$ $sk \leftarrow \mathcal{S}; pk \leftarrow \mu(sk); c^* \leftarrow \mathcal{C} \setminus \mathcal{V};$ $K_1^* \leftarrow \mathcal{D}_{sk}(c^*); K_0^* \leftarrow \mathcal{K}; b \leftarrow \{0, 1\};$ $b' \leftarrow \mathcal{A}^{\mathcal{O}}(\text{pub}, pk, c^*, K_b^*);$ If $b' = b$ then return 1 else return 0
	The definition of the oracle $\mathcal{O}$ in $\text{Expt}_{\Pi, \mathcal{A}}^{\text{CU}_2}$ and $\text{Expt}_{\Pi, \mathcal{A}}^{\text{CU}_1}$: $\mathcal{O}(c) = \begin{cases} \mathcal{D}_{sk}(c) & \text{If } c \in \mathcal{V} \\ \perp & \text{Otherwise} \end{cases}$

Fig. 3. The CU_2 experiment (left), the CU_1 experiment (upper-right), and the definition of the oracle (lower-right)

Priv: The (deterministic) private evaluation algorithm which takes $sk \in \mathcal{S}$ and an element $c \in \mathcal{C}$ as input, and outputs a hash value $K = \mathcal{D}_{sk}(c)$.

For all $\text{pub} \leftarrow \text{HGen}(1^k)$, we require the following: (1) for all $c \in \mathcal{C}$ and all $sk \in \mathcal{S}$, it holds that $\text{Priv}(sk, c) = \mathcal{D}_{sk}(c)$, and (2) for all $c \in \mathcal{V}$ with the corresponding witness w (about the fact that $c \in \mathcal{V}$), and all $sk \in \mathcal{S}$, it holds that $\text{Pub}(\mu(sk), c, w) = \text{Priv}(sk, c) = \mathcal{D}_{sk}(c)$.

Security Requirements. As usual, we define the subset membership problem for a HPS Π and its hardness.

Definition 4. We say that the subset membership problem in a HPS Π is hard if the following advantage function $\text{Adv}_{\Pi, \mathcal{A}}^{\text{SM}}(k)$ is negligible for any PPTA $\mathcal{A}$:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{SM}}(k) = |\Pr[\text{pub} \leftarrow \text{HGen}(1^k); b \leftarrow \{0, 1\}; c_1^* \leftarrow \mathcal{V}; c_0^* \leftarrow \mathcal{C} \setminus \mathcal{V}; \\ b' \leftarrow \mathcal{A}(\text{pub}, c_b^*) : b' = b] - \frac{1}{2}|$$

Hofheinz and Kiltz [17] defined the computational analogue of *strong universal₂* that is defined in [21] for a HPS, called *computational universal₂* (CU_2 security, for short), which we recall here. The CU_2 experiment $\text{Expt}_{\Pi, \mathcal{A}}^{\text{CU}_2}(k)$ for a HPS Π that an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ runs in is defined as in Fig. 3 (left). In the experiment, it is required that $\mathcal{A}_1$'s output c' satisfy $c' \in \mathcal{C} \setminus \mathcal{V}$ and $c' \neq c^*$.

Definition 5. We say that a HPS Π is CU_2 secure if the advantage function $\text{Adv}_{\Pi, \mathcal{A}}^{\text{CU}_2}(k) = |\Pr[\text{Expt}_{\Pi, \mathcal{A}}^{\text{CU}_2}(k) = 1] - 1/2|$ is negligible for any PPTA $\mathcal{A}$.

Moreover, we define a *universal₁*-analogue of CU_2 security, which we call *computational universal₁* (CU_1 security, for short). We define the CU_1 experiment $\text{Expt}_{\Pi, \mathcal{A}}^{\text{CU}_1}(k)$ for a HPS Π that an adversary $\mathcal{A}$ runs in as in Fig. 3 (upper-right).

Definition 6. We say that a HPS Π is CU_1 secure if the advantage function $\text{Adv}_{\Pi, \mathcal{A}}^{\text{CU}_1}(k) = |\Pr[\text{Expt}_{\Pi, \mathcal{A}}^{\text{CU}_1}(k) = 1] - 1/2|$ is negligible for any PPTA $\mathcal{A}$.

$\text{KG}(1^k) :$ $\text{pub} \leftarrow \text{HGen}(1^k)$ $sk \leftarrow \mathcal{S}; \quad pk \leftarrow \mu(sk)$ $PK \leftarrow (\text{pub}, pk)$ Return (PK, sk) .	$\text{Enc}(PK) :$ Pick $c \in \mathcal{V}$ uniformly together with a witness w . $K \leftarrow \text{Pub}(pk, c, w)$ Return (c, K) .	$\text{Dec}(sk, c) :$ $K \leftarrow \text{Priv}(sk, c)$ Return K .
---	---	--

Fig. 4. The KEM Γ_Π based on a HPS Π

Although CU_1 security is not explicitly defined in [17], it seems to us that this security is implicitly used for proving the CCCA security of the HPS-based KEM. Jiang and Wang [19] defined a slightly stronger version of CU_2 security which allows the second stage adversary $\mathcal{A}_2$ to have access to the oracle $\mathcal{O}$. This version of CU_2 security is satisfied by all known instantiations of HPS (see the following paragraph), and implies CU_1 security defined here. Thus, CU_1 security is not an additional security requirement for a HPS in practice. The reason why we introduce CU_1 security separately is that we believe that it makes our security analysis clearer. We also remark that CU_1 security is strictly weaker than “smoothness” defined in [19, Def. 7].

Concrete Instantiations of HPS. There are several known concrete instantiations of computational HPS that satisfy the above security requirements. The Kurosawa-Desmedt HPS [21,17] based on the DDH assumption, and its n -linear variant under the n -linear assumption [17], both of which are discrete logarithm-type constructions. Meanwhile, we also have a computational HPS based on the Paillier’s decision composite residuosity assumption [9]. For more details, see [9,17] and the references therein.

4.2 HPS-Based KEM and Bounded CCA Security

Let $\Pi = (\text{HGen}, \text{Pub}, \text{Priv})$ be a HPS. Then, the HPS-based KEM $\Gamma_\Pi = (\text{KG}, \text{Enc}, \text{Dec})$ [21,17] is constructed as in Fig. 4.

The following was shown by Hofheinz and Kiltz [17].

Theorem 5. ([17]) *If the subset membership problem of Π is hard, Π satisfies CU_2 and CU_1 security, then the HPS-based KEM Γ_Π is IND-CCCA secure.*

We show that under the same assumptions on the HPS used to prove its IND-CCCA security, the KEM Γ_Π satisfies IND-1-CCA security as well.

Theorem 6. *If the subset membership problem of Π is hard, Π satisfies CU_2 and CU_1 security, then the HPS-based KEM Γ_Π is IND-1-CCA secure.*

Intuition. CU_1 security of Π guarantees that, under the situation where the challenge ciphertext c^* is sampled from invalid elements (i.e. $c^* \leftarrow \mathcal{C} \setminus \mathcal{V}$), the real challenge session-key $K_1^* = \text{D}_{sk}(c^*)$ looks random to $\mathcal{A}$, as long as $\mathcal{A}$ ’s decapsulation query is a valid one. However, $\mathcal{A}$ is free to choose a ciphertext for a decapsulation query, and in particular, it can be invalid. This is the place where

CU₂ security comes into play. CU₂ security of Π guarantees that, even if $\mathcal{A}$'s decapsulation query c is an invalid one, $\mathcal{A}$ gets no significant information from the response, compared to just receiving a random value in $\mathcal{K}$, as long as $\mathcal{A}$ does not make any query after this query (and it is guaranteed because $\mathcal{A}$ is an IND-1-CCA adversary). Therefore, CU₁ and CU₂ together guarantee that $\mathcal{A}$'s decapsulation query essentially gives no significant information for distinguishing the real challenge session-key K_1^* from a random. Although CU₁ and CU₂ security are guaranteed only when the challenge ciphertext is an invalid one, the hardness of the subset membership problem in Π guarantees that $\mathcal{A}$'s behavior cannot be non-negligibly different between the case in which the challenge ciphertext is a valid one (i.e. in the original IND-1-CCA experiment), and the case in which the challenge ciphertext is an invalid one (and thus we can use CU₁ and CU₂ security of Π).

Proof Sketch. Let $\mathcal{A}$ be any PPTA IND-1-CCA adversary against the HPS-based KEM Γ_Π . Consider the following sequence of games.

Game 1. This is the original IND-1-CCA experiment, i.e. $\text{Expt}_{\Gamma_\Pi, \mathcal{A}}^{\text{IND-1-CCA}}(k)$.

Game 2. Same as Game 1, except that K_1^* is generated by $K_1^* \leftarrow \text{Priv}(sk, c^*)$.

Game 3. Same as Game 2, except that c^* is picked uniformly from $\mathcal{C} \setminus \mathcal{V}$.

Game 4. Same as Game 3, except that if $\mathcal{A}$'s decapsulation query c satisfies $c \in \mathcal{C} \setminus \mathcal{V}$, then it is answered with a uniformly random value $K \in \mathcal{K}$.

For $i \in [4]$, let S_i be the event that $\mathcal{A}$ succeeds in guessing the challenge bit (i.e. $b' = b$ occurs) in Game i . $\mathcal{A}$'s IND-1-CCA advantage can be estimated as follows:

$$\text{Adv}_{\Gamma_\Pi, \mathcal{A}}^{\text{IND-1-CCA}}(k) = |\Pr[S_1] - \frac{1}{2}| \leq \sum_{i \in [3]} |\Pr[S_i] - \Pr[S_{i+1}]| + |\Pr[S_4] - \frac{1}{2}|$$

The proof is completed by upperbounding each term in the right hand side of the above inequality to be negligible. We have $\Pr[S_1] = \Pr[S_2]$ due to the correctness of Π . $|\Pr[S_2] - \Pr[S_3]|$ is negligible due to the hardness of the subset membership problem in Π . $|\Pr[S_3] - \Pr[S_4]|$ and $|\Pr[S_4] - 1/2|$ can be shown to be negligible by using CU₂ security and CU₁ security of Π , respectively. $\square$

References

1. Abe, M., Gennaro, R., Kurosawa, K., Shoup, V.: Tag-KEM/DEM: A New Framework for Hybrid Encryption and A New Analysis of Kurosawa-Desmedt KEM. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 128–146. Springer, Heidelberg (2005)
2. Baek, J., Galindo, D., Susilo, W., Zhou, J.: Constructing Strong KEM from Weak KEM (or How to Revive the KEM/DEM Framework). In: Ostrovsky, R., De Prisco, R., Visconti, I. (eds.) SCN 2008. LNCS, vol. 5229, pp. 358–374. Springer, Heidelberg (2008)
3. Bellare, M., Desai, A., Pointcheval, D., Rogaway, P.: Relations among Notions of Security for Public-Key Encryption Schemes. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 26–45. Springer, Heidelberg (1998)

4. Bellare, M., Namprempre, C.: Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer, Heidelberg (2000)
5. Bellare, M., Sahai, A.: Non-malleable Encryption: Equivalence between Two Notions, and an Indistinguishability-Based Characterization. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 519–536. Springer, Heidelberg (1999); The revised version is available in Cryptology ePrint Archive (Report 2006/228)
6. Cash, D., Kiltz, E., Shoup, V.: The Twin Diffie-Hellman Problem and Applications. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 127–145. Springer, Heidelberg (2008)
7. Choi, S.G., Herranz, J., Hofheinz, D., Hwang, J.Y., Kiltz, E., Lee, D.H., Yung, M.: The Kurosawa-Desmedt key encapsulation is not chosen-ciphertext secure. Inf. Process. Lett. 109(16), 897–901 (2009)
8. Cramer, R., Hanaoka, G., Hofheinz, D., Imai, H., Kiltz, E., Pass, R., Shelat, A., Vaikuntanathan, V.: Bounded CCA2-Secure Encryption. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 502–518. Springer, Heidelberg (2007)
9. Cramer, R., Shoup, V.: Universal Hash Proofs and a Paradigm for Adaptive Chosen Ciphertext Secure Public-Key Encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002)
10. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM J. Computing 33(1), 167–226 (2003)
11. Dolev, D., Dwork, C., Naor, M.: Non-malleable cryptography. In: STOC 1991, pp. 542–552 (1991)
12. Hanaoka, G., Kurosawa, K.: Efficient Chosen Ciphertext Secure Public Key Encryption under the Computational Diffie-Hellman Assumption. In: Pieprzyk, J. (ed.) ASIACRYPT 2008. LNCS, vol. 5350, pp. 308–325. Springer, Heidelberg (2008)
13. Hanaoka, G., Kurosawa, K.: Between hashed DH and computational DH: Compact encryption from weaker assumption. IEICE Transactions E93-A(11), 1994–2006 (2010)
14. Haralambiev, K., Jager, T., Kiltz, E., Shoup, V.: Simple and Efficient Public-Key Encryption from Computational Diffie-Hellman in the Standard Model. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 1–18. Springer, Heidelberg (2010)
15. Herranz, J., Hofheinz, D., Kiltz, E.: Some (in)sufficient conditions for secure hybrid encryption. Inf. Comput. 208(11), 1243–1257 (2010)
16. Hiwatari, H., Tanaka, K., Asano, T., Sakumoto, K.: Multi-recipient Public-Key Encryption from Simulators in Security Proofs. In: Boyd, C., González Nieto, J. (eds.) ACISP 2009. LNCS, vol. 5594, pp. 293–308. Springer, Heidelberg (2009)
17. Hofheinz, D., Kiltz, E.: Secure Hybrid Encryption from Weakened Key Encapsulation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 553–571. Springer, Heidelberg (2007)
18. Hohenberger, S., Lewko, A., Waters, B.: Detecting Dangerous Queries: A New Approach for Chosen Ciphertext Security. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 663–681. Springer, Heidelberg (2012); The full version is available in Cryptology ePrint Archive (Report 2012/006)
19. Jiang, S., Wang, H.: Plaintext-Awareness of Hybrid Encryption. In: Pieprzyk, J. (ed.) CT-RSA 2010. LNCS, vol. 5985, pp. 57–72. Springer, Heidelberg (2010)

20. Kiltz, E., Pietrzak, K., Stam, M., Yung, M.: A New Randomness Extraction Paradigm for Hybrid Encryption. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 590–609. Springer, Heidelberg (2009)
21. Kurosawa, K., Desmedt, Y.: A New Paradigm of Hybrid Encryption Scheme. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 426–442. Springer, Heidelberg (2004)
22. Matsuda, T., Matsuura, K.: Parallel Decryption Queries in Bounded Chosen Ciphertext Attacks. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) PKC 2011. LNCS, vol. 6571, pp. 246–264. Springer, Heidelberg (2011)
23. Nagao, W., Manabe, Y., Okamoto, T.: On the equivalence of several security notions of KEM and DEM. IEICE Transactions E91-A(1), 283–297 (2008)
24. Pass, R., Shelat, A., Vaikuntanathan, V.: Relations Among Notions of Non-malleability for Encryption. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 519–535. Springer, Heidelberg (2007)
25. Phan, D.H., Pointcheval, D.: About the Security of Ciphers (Semantic Security and Pseudo-Random Permutations). In: Handschuh, H., Hasan, M.A. (eds.) SAC 2004. LNCS, vol. 3357, pp. 182–197. Springer, Heidelberg (2004)
26. Rackoff, C., Simon, D.R.: Non-interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 433–444. Springer, Heidelberg (1992)
27. Sakai, Y., Hanaoka, G., Kurosawa, K., Ohta, K.: A Generic Method for Reducing Ciphertext Length of Reproducible KEMs in the RO Model. In: Echizen, I., Kunihiro, N., Sasaki, R. (eds.) IWSEC 2010. LNCS, vol. 6434, pp. 55–69. Springer, Heidelberg (2010)