

Subtleties in the Definition of IND-CCA: When and How Should Challenge-Decryption be Disallowed?

MIHIR BELLARE¹

DENNIS HOFHEINZ²

EIKE KILTZ³

Abstract

The definition of IND-CCA disallows an adversary from querying the challenge ciphertext to its decryption oracle. We point out that there are several ways to formalize this. We show that, surprisingly, for public-key encryption the resulting notions are not all equivalent. We then consider the same question for key-encapsulation mechanisms (KEMs) and show that in this case the four notions *are* all equivalent. Our discoveries are another manifestation of the subtleties that make the study of cryptography so attractive and are important towards achieving the definitional clarity and unity required for firm foundations.

Keywords: Definitions, foundations, encryption, chosen-ciphertext attack

1 Introduction

Cryptography is founded on definitions. Results in cryptography are meaningful, clear or useful to the extent that this is true of the definitions they target. To this already strong impetus for the careful study of definitions one must add the sheer intellectual satisfaction that arises from the discovery of the hidden depths of the notions involved.

This paper strengthens the foundations of the theory of encryption by identifying a subtlety in the definition of IND-CCA and studying the relations between the variant notions that arise.

1.1 The PKE case

We begin by recalling the definitional template. The underlying experiment picks a public key pk and matching secret key sk , and then provides pk to the adversary A . The latter runs in two phases in *both* of which A has access to an oracle for decryption under sk . It ends its first phase by outputting a pair M_0, M_1 of messages. The experiment picks a challenge bit b at random, encrypts M_b under pk , and returns the resulting *challenge ciphertext* C^* to A . The latter now enters its second phase, which it ends by outputting a bit b' . We say that A wins if $b = b'$. Security requires that the probability of winning minus $1/2$ is negligible.

If A can query the challenge ciphertext C^* to its decryption oracle, it can easily win the above game. The definition accordingly disallows such a *challenge decryption* query.

At first glance this “no-challenge-decryption” condition seems clear and unambiguous. A closer look shows otherwise. We now discuss two issues or dimensions in the formalization and see how this gives rise to four possible notions of IND-CCA that we will relate.

It is clear that we must disallow a challenge decryption query in the second phase of the attack, but what about the first? To be more precise, let S_j denote the set of all decryption queries made by A in phase j ($j = 1, 2$). Then we have two options: at the end of the experiment, when we can evaluate this condition, either disallow $C^* \in S_2$ (denote this “S” for “second”) or disallow $C^* \in S_1 \cup S_2$ (denote this “B” for “both”). The basic rationale for the no-challenge-decryption condition, namely that if the

¹Dept. of Computer Science & Engineering, University of California San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. Email: mihir@cs.ucsd.edu. URL: <http://www.cs.ucsd.edu/users/mihir>.

²CWI Amsterdam, The Netherlands. Email: hofheinz@cwi.nl. URL: <http://www.cwi.nl/~hofheinz/>.

³CWI Amsterdam, The Netherlands. Email: kiltz@cwi.nl. URL: <http://kiltz.net>.

	A wins if	A is valid if
IND-CCA-SP	$(b = b') \wedge (C^* \notin S_2)$	
IND-CCA-BP	$(b = b') \wedge (C^* \notin S_1 \cup S_2)$	
IND-CCA-SE	$(b = b')$	$(C^* \notin S_2)$
IND-CCA-BE	$(b = b')$	$(C^* \notin S_1 \cup S_2)$

Figure 1: Summary of our IND-CCA notions for PKE.

adversary queries C^* it wins trivially, holds true regardless of the phase in which the query is made and thus supports either choice.

The existence of this choice having been pointed out, one’s first reaction may be that it does not matter, meaning the two are equivalent. This turns out not to be true. Before we get there, however, let us discuss another definitional issue. Namely, what exactly does “disallow” mean? Again there are two options. The first option is to have the experiment, after the adversary has completed, test whether C^* is in an undesired set (S_2 or $S_1 \cup S_2$, depending on whether we do “S” or “B”) and, if so, return false, meaning declaring the adversary to have lost. We call this a *penalty* (“P”) style notion since the adversary is being penalized, a posteriori, for its actions. In the literature however it is more common to not have the experiment impose a penalty but just say, outside of the experiment, that the adversary is “not allowed” or just “may not” make a challenge decryption query. But what exactly (meaning, formally) does this mean? It seems to us that the natural interpretation, and the one intended by the authors, is that we are quantifying over all (polynomial-time) adversaries that *never* make a challenge decryption query, meaning have zero probability of doing so in the experiment. We refer to this as an *exclusion* (“E”) style notion since certain adversaries are a priori excluded from consideration.

With two options (“B” or “S”) in the first dimension and another two (“P” or “E”) in the second we obtain four notions. Figure 1 summarizes them. The first column shows the winning condition for A, namely, the condition under which the experiment returns true. The second column shows when A is valid, meaning we quantify only over (polynomial-time) adversaries for which the validity condition holds with probability one in the experiment. See Section 3 for formal definitions.

The left-hand side of Figure 2 summarizes the relations we show between the notions. An implication $\text{IND-CCA-X} \rightarrow \text{IND-CCA-Y}$ means every PKE scheme that is IND-CCA-X secure is also IND-CCA-Y secure. A separation $\text{IND-CCA-X} \not\rightarrow \text{IND-CCA-Y}$ means we give an example of a PKE scheme that is IND-CCA-X secure but not IND-CCA-Y secure. Only a minimal set of relations is explicitly shown; others follow. For example, $\text{IND-CCA-BE} \not\rightarrow \text{IND-CCA-SE}$, since otherwise we would contradict shown separations.

These results show that disallowing a challenge-decryption query in both phases results in a strictly weaker notion than disallowing it only in the second phase, and this is true for both penalty and exclusion style formulations. That is, IND-CCA-SP and IND-CCA-BP are not equivalent, and also IND-CCA-SE and IND-CCA-BE are not equivalent. Another interesting fact is that if the challenge decryption query is disallowed only in the second phase then it makes no difference whether this is by penalty or exclusion (that is, IND-CCA-SE and IND-CCA-SP are equivalent), but, in contrast if the challenge decryption query is disallowed in both phases, an exclusion style formulation results in a strictly weaker notion than a penalty style formulation (that is, IND-CCA-BE does not imply IND-CCA-BP). One of the conclusions from this is that the “S” notions should be preferred, not only because they are stronger but also because the penalty and exclusion style formulations are equivalent.

One might at first think that (contrary to our claim) IND-CCA-SP and IND-CCA-BP are equivalent. Why? To explain, let us say that a PKE scheme is “smooth” if the number of possible ciphertexts is large (super-polynomial) for any message. (See Appendix A for a more precise definition.) Now reason as follows: first, any smooth IND-CCA-BP scheme is IND-CCA-SP since the adversary cannot predict,

hence query, the challenge ciphertext in the first phase; second, even an IND-CPA scheme must be smooth, else we could break it by re-encrypting the challenge messages until the challenge ciphertext is seen. What’s the catch? It is that the second claim is false. As our proof of Theorem 3.1 shows, even an IND-CCA-BP (let alone IND-CPA) scheme need not be smooth: “weak” messages, meaning ones with few corresponding ciphertexts, can exist without contradicting IND-CCA-BP security as long as they are hard to find without access to a decryption oracle.¹

Our work was sparked by seeing variations in the formalization of the “no-challenge-decryption” condition in the literature. For example, [4, 13, 19, 28, 39, 29, 37] define what in our taxonomy is IND-CCA-SE. However, many works [11, 12, 20, 32, 33, 34, 41] simply have a phrase like “the adversary is not allowed to query the challenge ciphertext to the decryption oracle.” On the one hand, since no phase is indicated, this could be interpreted as IND-CCA-BE. On the other hand, since the challenge ciphertext is not defined in the first phase, it could be interpreted as IND-CCA-SE. But our results say that these notions are different.

Penalty-style formulations are rarer, but [2] defines IND-CCA-SP and [1] defines IND-CCA-BP. (This definition is for HIBEs, but this gives PKE for hierarchies of depth 0.) The single-user definition in [3] is IND-CCA-SE but the multi-user definition is in the BE style. Moving to textbooks, Goldreich [21, Sec 5.4.1.1], Delfs and Knebel [15, Def 9.17] and Katz and Lindell [24, Sec 10.6] define IND-CCA-SE while Menezes, Van Oorschot and Vanstone [30, Sec 8.1.1] define IND-CCA-BE.

In order to have firm foundations—in particular a unique interpretation and common understanding of results—it is important to have definitional unity, meaning that different definitions intending or claiming to represent the same notion should really do so. Our work is a step to this end. Our work also highlights a general definitional issue that we feel needs to be addressed with more care. Namely, in many instances one has a choice between formalizing something in a penalty or exclusion style. One should take care to ascertain that the resulting notions are equivalent, for as our results show this is not always true. Finally, we think our results are an interesting illustration of how seemingly minor definitional elements affect the power of the notion.

1.2 The KEM case

Cramer and Shoup [14] show that an IND-CCA PKE scheme can be obtained by combining an IND-CCA KEM (Key Encapsulation Mechanism) with an IND-CCA DEM (Data Encapsulation Algorithm). This has proved to be a powerful and useful paradigm, leading to increased interest in KEMs [8, 16, 25, 26, 40]. When, in this light, we revisit the definition of IND-CCA for KEMs we find that there arise the same issues regarding challenge decryption as in the PKE case. We again obtain four notions that we denote as before, with the notion of [14], in our taxonomy, being IND-CCA-SE. Our results resolving the relations among the notions are depicted on the right-hand side of Figure 2. We see an interesting contrast with the PKE case of the left side of the same figure, namely that in the KEM case the notions are all equivalent. Intuitively this is true because in the KEM case the role of the encrypted “message” is played by a symmetric key not under adversarial control. Our results make crucial use of smoothness: we show that IND-CCA-BP implies IND-CCA-SP (unlike for PKE) by first showing that any smooth IND-CCA-BP KEM is smooth (unlike for PKE) and then showing that any smooth IND-CCA-BP KEM is IND-CCA-SP (this was true also for PKE).

In addition we show that both the penalty and exclusion versions (IND-CCA-OP and IND-CCA-OE) of a simple one-phase definition of IND-CCA for KEMs are equivalent to all the others, simplifying the task of showing that specific KEMs are IND-CCA secure. IND-CCA-OE was proposed by [26] who showed it is equivalent to IND-CCA-SP when the KEM encapsulation algorithm induces a uniform distribution

¹The first claim above—namely that IND-CCA-BP implies IND-CCA-SP for smooth schemes—is actually true, and useful because “real” schemes are typically (unconditionally) smooth. Interestingly, IND-CCA-BE fails to imply IND-CCA-SE even for smooth schemes, indicating a further weakness of exclusion-style formulations. See Appendix A for more information.




Figure 2: Relations between the various IND-CCA security notions for PKE schemes (left) and KEMs (right). An arrow $\text{IND-CCA-X} \rightarrow \text{IND-CCA-Y}$ is an implication and a barred arrow $\text{IND-CCA-X} \not\rightarrow \text{IND-CCA-Y}$ is a separation. Dotted lines denote trivial implications. The numbers next to the solid lines indicate the theorems establishing them.

on the keyspace, an assumption we don't make.

1.3 Extensions and related work

Our results for PKE extend also to private-key (i.e. symmetric) encryption, IBE (Identity-Based Encryption) and HIBE (Hierarchical IBE). That is, the same four notions again emerge and the relations are as shown on the left-hand-side of Figure 2. In the (H)IBE case, most works [7, 27] define IND-CCA-SE but [6] defines IND-CCA-BE.

The notion of Naor and Yung [31] gives the adversary the decryption oracle *only* in the first phase. This is sometimes called a non-adaptive attack. When we talk of IND-CCA in this paper, we mean under adaptive attack: all our notions give the adversary the decryption oracle in both phases. This form of IND-CCA is often attributed to Rackoff and Simon [36]. They were indeed the first to consider adaptive attacks, but they give the adversary access to the decryption oracle *only* in the second phase—which, as shown by [34], is strictly weaker than giving access in both phases—and their definition is only for random one bit messages. Dolev, Dwork and Naor [17] do not formally define IND-CCA but their definition of non-malleability under CCA selects the “SE” option. Definitions of IND-CCA of the form that is now common seem to begin with the concurrent 1998 works [4, 13].

In the context of relaxed CCA security (RCCA security, [10, 22, 35]), a variant of the IND-CCA-SE definition is employed. In the RCCA definition, the adversary gets a completely unrestricted decryption oracle in the first phase. In the second phase, the adversary may ask for arbitrary decryptions. *However*, if the decrypted message equals is one of the two adversarially chosen challenge messages m_0, m_1 , then the the adversary simply gets a special answer “test” (or “invalid” in [22]) that indicates that either m_0 or m_1 is the plaintext. (This rule applies in particular to a decryption of the challenge ciphertext.)

We stress that the RCCA security constitutes a weakening of the IND-CCA-SE definition that is orthogonal to our notion of IND-CCA-BE. In particular, we consider different formalizations that reflect the same intuitive definition (security under unrestricted chosen-ciphertext attacks), while RCCA security captures a different intuition (re-randomizing the challenge ciphertext is explicitly allowed).

The RCCA and IND-CCA security notions have been proven equivalent to realizing ideal functionalities in the framework of Universal Composability [9]. In these proofs [10, 23], the IND-CCA-SE variant of IND-CCA security was used. This is another a hint that the “S” notions are the “right” notions to use.

2 Preliminaries

If x is a string, then $|x|$ denotes its length, while if S is a set then $|S|$ denotes its size. If $k \in \mathbb{N}$ then 1^k denotes the string of k ones. If S is a set then $s \leftarrow_{\mathbb{R}} S$ denotes the operation of picking an element s of S uniformly at random. Unless otherwise indicated, algorithms are randomized and polynomial time.

<p>Experiment $\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-X}}(k)$</p> <p>$(pk, sk) \leftarrow_{\text{R}} \text{Kg}(1^k); S_1, S_2 \leftarrow \emptyset$</p> <p>$(M_0, M_1, St) \leftarrow_{\text{R}} A_1^{\text{DEC}_1(\cdot)}(1^k, pk)$</p> <p>$b \leftarrow_{\text{R}} \{0, 1\}; C^* \leftarrow_{\text{R}} \text{Enc}(pk, M_b)$</p> <p>$b' \leftarrow_{\text{R}} A_2^{\text{DEC}_2(\cdot)}(C^*, St)$</p> <p>Return:</p> <p>SE, BE : $(b = b')$</p> <p>SP : $(b = b') \wedge (C^* \notin S_2)$</p> <p>BP : $(b = b') \wedge (C^* \notin S_1 \cup S_2)$</p>	<p>Oracle $\text{DEC}_1(C)$</p> <p>$S_1 \leftarrow S_1 \cup \{C\}$</p> <p>return $\text{Dec}(sk, C)$</p> <p>Oracle $\text{DEC}_2(C)$</p> <p>$S_2 \leftarrow S_2 \cup \{C\}$</p> <p>return $\text{Dec}(sk, C)$</p>
--	---

Figure 3: Experiment $\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-X}}(k)$ for $X \in \{\text{SE}, \text{BE}, \text{SP}, \text{BP}\}$. The experiments differ only in how they compute their final Boolean output, which depends on X as shown.

By $z \leftarrow_{\text{R}} A^{\mathcal{O}_1, \mathcal{O}_2, \dots}(x, y, \dots)$ we denote the operation of running algorithm A with inputs x, y, \dots and access to oracles $\mathcal{O}_1, \mathcal{O}_2, \dots$, and letting z be the output. An adversary is an algorithm or a tuple of algorithms.

The advantage of an adversary I in inverting a function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is defined for $k \in \mathbb{N}$ as

$$\mathbf{Adv}_{f,I}^{\text{ow}}(k) = \Pr[f(x) = f(y) : x \leftarrow_{\text{R}} \{0, 1\}^k; y \leftarrow_{\text{R}} I(1^k, f(x))].$$

We say that f is one-way if $\mathbf{Adv}_{f,I}^{\text{ow}}(\cdot)$ is negligible for all adversaries I . We say that f is a permutation if for all $k \in \mathbb{N}$, the restriction of f to $\{0, 1\}^k$ is a permutation on $\{0, 1\}^k$.

3 Results for Public-Key Encryption

We begin with definitions.

SYNTAX. An asymmetric encryption scheme $\text{PKE} = (\text{Kg}, \text{Enc}, \text{Dec})$ is a triple of algorithms. The key generation algorithm Kg takes a security parameter 1^k and returns a pair (pk, sk) of matching public and secret keys. The encryption algorithm Enc takes a public key pk and a message $M \in \{0, 1\}^*$ to produce a ciphertext C . The deterministic decryption algorithm Dec takes sk and ciphertext C to produce either a message $M \in \{0, 1\}^*$ or a special symbol \perp to indicate that the ciphertext was invalid. The consistency requirement is that for all $k \in \mathbb{N}$, for all (pk, sk) which can be output by $\text{Kg}(1^k)$, for all $M \in \{0, 1\}^*$, and for all C that can be output by $\text{Enc}(pk, M)$, we have that $\text{Dec}(sk, C) = M$.

IND-CCA SECURITY. We first provide formal definitions and then explanations. An IND-CCA adversary $A = (A_1, A_2)$ is a pair of algorithms such that the output of A_1 is always a tuple (M_0, M_1, St) satisfying $|M_0| = |M_1|$. Let \mathcal{A} be the class of all such adversaries. Let $X \in \{\text{SP}, \text{BP}, \text{SE}, \text{BE}\}$. To an adversary $A = (A_1, A_2) \in \mathcal{A}$, a PKE scheme $\text{PKE} = (\text{Kg}, \text{Enc}, \text{Dec})$ and $k \in \mathbb{N}$, we associate the experiment $\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-X}}(k)$ of Figure 3. We define the advantage of A as

$$\mathbf{Adv}_{\text{PKE},A}^{\text{ind-cca-X}}(k) = 2 \Pr[\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-X}}(k) \Rightarrow \text{true}] - 1.$$

Let $\mathcal{A}_{\text{PKE}}^{\text{SP}} = \mathcal{A}_{\text{PKE}}^{\text{BP}} = \mathcal{A}$ be the class of all IND-CCA adversaries. Let $\mathcal{A}_{\text{PKE}}^{\text{SE}}$ be the class of all $A \in \mathcal{A}$ such that for all $k \in \mathbb{N}$, the probability that $C^* \in S_2$ in $\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-SE}}(k)$ is 0. Let $\mathcal{A}_{\text{PKE}}^{\text{BE}}$ be the class of all $A \in \mathcal{A}$ such that for all $k \in \mathbb{N}$, the probability that $C^* \in S_1 \cup S_2$ in $\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-BE}}(k)$ is 0. We say that PKE is IND-CCA- X secure if $\mathbf{Adv}_{\text{PKE},A}^{\text{ind-cca-X}}(\cdot)$ is negligible for all $A \in \mathcal{A}_{\text{PKE}}^X$.

DISCUSSION. These notions reflect the different treatments of challenge decryption queries along two dimensions. The first dimension is whether decryption of the challenge ciphertext is disallowed in both

(“B”) phases or only in the second (“S”) phase. The second dimension is how, technically, to disallow this query. Here the first choice is that the experiment penalizes (“P”) the adversary by returning “false” if it makes a disallowed query, and the second choice (“E”) is that adversaries with non-zero probability of making the disallowed query are simply not considered.

There is another option in the second dimension, namely that one considers the class of adversaries that have negligible (rather than zero) probability of making a query of the unallowed type. We do not consider this since we have not found it defined or indicated in the literature. Indeed, the intent of a typical phrase of the form “the adversary is not allowed to query the challenge ciphertext to the decryption oracle” seems to be that such a query is *never* allowed. Had the writers meant allowed only with negligible probability, one would have expected it precisely stated as such.

Trivial implications. The trivial implications (dashed arrows) from Figure 2 should be clear from the definitions. Briefly, IND-CCA-SP implies IND-CCA-SE because if the probability that $C^* \in S_2$ is zero then the winning conditions $(b = b')$ and $(b = b') \wedge (C^* \in S_2)$ are equivalent. The reason for IND-CCA-BP implying IND-CCA-BE is analogous. IND-CCA-SP implies IND-CCA-BP because the winning condition of the latter is more stringent than that of the former. IND-CCA-SE implies IND-CCA-BE because $\mathcal{A}_{\text{PKE}}^{\text{BE}} \subseteq \mathcal{A}_{\text{PKE}}^{\text{SE}}$.

IND-CCA-BP $\not\Rightarrow$ IND-CCA-SP. Theorem 3.1 below shows that for penalty-style notions, disallowing a challenge-ciphertext query in both phases results in a notion strictly weaker than that resulting from disallowing it only in the second phase. That this is also true for the exclusion-style notions will follow by combining Theorems 3.1 and 3.2.

Theorem 3.1 [IND-CCA-BP $\not\Rightarrow$ IND-CCA-SP] Assume there exist one-way permutations and a scheme PKE which is IND-CCA-BP secure. Then there exists a scheme $\overline{\text{PKE}}$ which is IND-CCA-BP secure but not IND-CCA-SP secure. ■

Proof: We want to design a scheme $\overline{\text{PKE}} = (\overline{\text{Kg}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ which is IND-CCA-BP secure but not IND-CCA-SP secure. That is, ability to query the challenge ciphertext in the first phase should lead to an attack, but, when this is disallowed, the scheme should be secure. The intuition is as follows. Suppose there was a special message M_{weak} and a special ciphertext C_{weak} such that $\overline{\text{Enc}}(pk, M_{\text{weak}})$ always (meaning, with probability one) returns C_{weak} . Then an adversary could output as its challenge messages $M_0 = M_{\text{weak}}$ and some $M_1 \neq M_{\text{weak}}$. If the challenge bit is 0 then the challenge ciphertext C^* must be C_{weak} , and otherwise (by consistency) must be different from C_{weak} , so, given C^* the adversary can always determine the challenge bit, and the scheme is not IND-CCA-SP. The difficulty is that it is not IND-CCA-BP either. (In fact, it is not even IND-CPA.) To make it IND-CCA-BP, we ensure that M_{weak} can only be found by querying C_{weak} to the decryption oracle in the first phase. However, there is now a difficulty. Namely, the encryption algorithm $\overline{\text{Enc}}$ needs to return C_{weak} given pk, M_{weak} , meaning must at some level know M_{weak} . Yet the adversary, who is given pk, C_{weak} , and the description of $\overline{\text{Enc}}$, must not know M_{weak} . (Unless it queries C_{weak} to the decryption oracle.) To ensure this, we put in pk an image of M_{weak} under a one-way permutation. Then neither pk nor $\overline{\text{Enc}}$ reveal M_{weak} , but $\overline{\text{Enc}}$ can test whether a given input equals M_{weak} . We now proceed to the details.

Let $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be a one-way permutation and assume that $\text{PKE} = (\text{Kg}, \text{Enc}, \text{Dec})$ is IND-CCA-BP secure. Consider the scheme $\overline{\text{PKE}} = (\overline{\text{Kg}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ whose constituent algorithms are shown in Figure 4, where N_k is set to $\{1^k\}$. The ciphertext C_{weak} from the above discussion is $(1, 1^k)$. Now we want to claim that $\overline{\text{PKE}}$ is IND-CCA-BP secure but not IND-CCA-SP secure. However, we first check that $\overline{\text{PKE}}$ is consistent. The reason we want to highlight this (usually trivial) check is that it is the (only) place we use the assumption that f is a permutation rather than just a function.

Claim 1. $\overline{\text{PKE}}$ is consistent.

Proof. We have to show that $\overline{\text{Dec}}(\overline{sk}, \overline{\text{Enc}}(\overline{pk}, M)) = M$, always. If $f(M) \neq Y$ where $\overline{pk} = (pk, Y)$, this follows from the consistency of PKE. So suppose $f(M) = Y$. In that case $\overline{\text{Enc}}(pk, M)$ returns $\overline{C} = (1, 1^k)$

Alg $\overline{\text{Kg}}(1^k)$ $(pk, sk) \leftarrow_{\text{R}} \text{Kg}(k)$ $M_{\text{weak}} \leftarrow_{\text{R}} \{0, 1\}^k$ $Y \leftarrow f(M_{\text{weak}})$ $\overline{pk} \leftarrow (pk, Y)$ $\overline{sk} \leftarrow (sk, M_{\text{weak}})$ Return $(\overline{pk}, \overline{sk})$	Alg $\overline{\text{Enc}}(\overline{pk}, M)$ Parse $(pk, Y) \leftarrow \overline{pk}$ If $f(M) = Y$ then $w \leftarrow_{\text{R}} N_k$ $\overline{C} \leftarrow (1, w)$ Else $C \leftarrow_{\text{R}} \text{Enc}(pk, M)$ $\overline{C} \leftarrow (0, C)$ Return \overline{C}	Alg $\overline{\text{Dec}}(\overline{sk}, \overline{C})$ Parse $(sk, M_{\text{weak}}) \leftarrow \overline{sk}$ Parse $(s, C) \leftarrow \overline{C}$ If $s = 0$ then return $\text{Dec}(sk, C)$ If $s = 1$ and $C \in N_k$ then $\text{return } M_{\text{weak}}$ Return \perp
--	---	---

Figure 4: Counterexample scheme $\overline{\text{PKE}}$ for proofs of Theorems 3.1 and 3.3. In the first case $N_k = \{1^k\}$ and in the second case $N_k = \{0, 1\}^k$.

Alg $A_1^{\text{Dec}_1(\cdot)}(pk)$ $bad \leftarrow \text{false}; T_1, T_2, D_1, D_2 \leftarrow \emptyset$ $M_{\text{weak}} \leftarrow_{\text{R}} \{0, 1\}^k; Y \leftarrow f(M_{\text{weak}})$ $\overline{pk} \leftarrow (pk, Y)$ $(M_0, M_1, \overline{St}) \leftarrow_{\text{R}} B_1^{\overline{\text{SDEC}}_1(\cdot)}(pk)$ If $M_{\text{weak}} \in \{M_0, M_1\}$ and $(1, 1^k) \in T_1$ then $bad \leftarrow \text{true}$ $M_0, M_1 \leftarrow_{\text{R}} \{0, 1\}^k \setminus D_1$ $St \leftarrow (\overline{St}, bad)$ Return (M_0, M_1, St)	Alg $A_2^{\text{DEC}_2(\cdot)}(C^*, St)$ Parse $(\overline{St}, bad) \leftarrow St$ If $bad = \text{true}$ then $b' \leftarrow_{\text{R}} \{0, 1\}$ $\text{return } b'$ $\overline{C}^* \leftarrow (0, C^*)$ $b' \leftarrow_{\text{R}} B_2^{\overline{\text{SDEC}}_2(\cdot)}(\overline{C}^*, \overline{St})$ Return b'	Sub $\overline{\text{SDEC}}_j(\overline{C})$ $T_j \leftarrow T_j \cup \{\overline{C}\}$ Parse $(s, C) \leftarrow \overline{C}$ If $s = 0$ then $M \leftarrow \text{DEC}_j(C)$ $D_j \leftarrow D_j \cup \{M\}$ $\text{return } M$ If $(s, C) = (1, 1^k)$ then $D_j \leftarrow D_j \cup \{M_{\text{weak}}\}$ $\text{return } M_{\text{weak}}$ Return \perp
---	--	---

Figure 5: Adversary $A = (A_1, A_2) \in \mathcal{A}_{\overline{\text{PKE}}}^{\text{BP}}$ for the proof of Claim 3.

which is decrypted by $\overline{\text{Dec}}$ to M_{weak} . Since f is a permutation (hence in particular injective) we have that $M_{\text{weak}} = M$. \square

Claim 2. $\overline{\text{PKE}}$ is not IND-CCA-SP secure.

Proof. Consider adversary $A = (A_1, A_2) \in \mathcal{A}_{\overline{\text{PKE}}}^{\text{SP}}$ that proceeds as follows. Given $\overline{pk} = (pk, Y)$, algorithm A_1 queries $\text{DEC}_1(\cdot)$ on ciphertext $C = (1, 1^k)$ to obtain M_{weak} . It picks $M_1 \leftarrow_{\text{R}} \{0, 1\}^k \setminus \{M_{\text{weak}}\}$ and returns $M_0 = M_{\text{weak}}$ and M_1 as the two challenge messages. A_2 obtains a challenge ciphertext C^* and returns $b' = 0$ if $C^* = (1, 1^k)$ and $b' = 1$, otherwise. We have $\text{Adv}_{\overline{\text{PKE}}, A}^{\text{ind-cca-SP}}(k) = 1$. Note that with probability $1/2$, A queries the challenge ciphertext to the decryption oracle in the first phase which is why this does not show $\overline{\text{PKE}}$ is IND-CCA-BP insecure. \square

Claim 3. $\overline{\text{PKE}}$ is IND-CCA-BP secure.

Proof. Given an adversary $B = (B_1, B_2) \in \mathcal{A}_{\overline{\text{PKE}}}^{\text{BP}}$ we build $A = (A_1, A_2) \in \mathcal{A}_{\overline{\text{PKE}}}^{\text{BP}}$ and an adversary I against the one way permutation f such that, for all $k \in \mathbb{N}$,

$$\text{Adv}_{\overline{\text{PKE}}, B}^{\text{ind-cca-BP}}(k) \leq \text{Adv}_{\overline{\text{PKE}}, A}^{\text{ind-cca-BP}}(k) + 2\text{Adv}_{f, I}^{\text{ow}}(k). \quad (1)$$

We start by describing $A = (A_1, A_2)$ in Figure 5. Here, A simulates the oracles of B using the shown subroutines $\overline{\text{SDEC}}_j(\cdot)$ ($j = 1, 2$). For B , this provides a perfect simulation of experiment $\text{Exp}_{\overline{\text{PKE}}, B}^{\text{ind-cca-BP}}$ unless $M_{\text{weak}} \in \{M_0, M_1\}$. This motivates the definition of the following events. Event BD is that

$M_{\text{weak}} \in \{M_0, M_1\}$. Event ASK is that B_1 asks for the decryption of $\overline{C} = (1, 1^k)$. We have

$$\begin{aligned} & \Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \right] \\ &= \Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \neg \text{BD} \right] + \Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \text{BD} \right]. \end{aligned} \quad (2)$$

The following takes care of the first summand and uses that A provides a good view for B unless BD occurs, and that the probability for BD is the same in both experiments:

$$\Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \neg \text{BD} \right] = \Pr \left[\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \neg \text{BD} \right]. \quad (3)$$

To bound the second summand of (2), we start with

$$\begin{aligned} & \Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \text{BD} \right] \\ & \leq \Pr [\text{BD} \wedge \neg \text{ASK}] + \Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \text{BD} \wedge \text{ASK} \right]. \end{aligned} \quad (4)$$

We design an adversary I against the one-wayness of f such that

$$\Pr [\text{BD} \wedge \neg \text{ASK}] \leq \mathbf{Adv}_{f,I}^{\text{ow}}(k). \quad (5)$$

I gets $Y = f(M_{\text{weak}})$ for uniformly chosen $M_{\text{weak}} \in \{0, 1\}^k$ and tries to compute M_{weak} . To this end, I proceeds as follows:

<p>Alg I(Y)</p> <p>$(pk, sk) \leftarrow_{\text{R}} \text{Kg}(1^k); \overline{pk} \leftarrow (pk, Y)$</p> <p>$(M_0, M_1, St) \leftarrow_{\text{R}} B_1^{\text{SDEC}_1(\cdot)}(1^k, \overline{pk})$</p> <p>If $f(M_0) = Y$ then return M_0</p> <p>If $f(M_1) = Y$ then return M_1</p> <p>Else return \perp</p>	<p>Oracle $\overline{\text{SDEC}}_1(\overline{C})$</p> <p>Parse $(b, C) \leftarrow \overline{C}$</p> <p>If $b = 0$ then return $\text{Dec}(sk, C)$</p> <p>Else return \perp</p>
---	---

Note that B_1 has exactly the same view in experiment $\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}$ and in the simulation inside I *unless* it asks for a decryption of $(1, 1^k)$. Also, I is successful in inverting f iff $M_{\text{weak}} \in \{M_0, M_1\}$. Hence, Equation (5) is true.

Note that the probability of $\text{BD} \wedge \text{ASK}$ could be high, because nothing prevents B_1 from making the decryption query $(1, 1^k)$ to get M_{weak} and then setting either M_0 or M_1 to M_{weak} . However, we note that if $\text{BD} \wedge \text{ASK}$ does occur, then B loses with probability 1/2 because $\overline{C}^* = (1, 1^k)$ with that probability. That is,

$$\Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \mid \text{BD} \wedge \text{ASK} \right] \leq 1/2 \quad (6)$$

On the other hand,

$$\Pr \left[\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \mid \text{BD} \wedge \text{ASK} \right] = 1/2. \quad (7)$$

This is because if $\text{BD} \wedge \text{ASK}$ happens then A_1 sets *bad* to **true** and A_2 returns a random decision b' . Here we also use that by consistency of the scheme, picking M_0, M_1 from $\{0, 1\}^k \setminus D_1$, ensures that A_1 never queries the challenge ciphertext to the decryption oracle in the first phase. Now note that the probability of $\text{BD} \wedge \text{ASK}$ is the same in both experiments. Hence, from (6),(7), we get

$$\Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \text{BD} \wedge \text{ASK} \right] \leq \Pr \left[\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \text{BD} \wedge \text{ASK} \right].$$

Combining this with (4) and (5) yields

$$\Pr \left[\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \text{BD} \right] \leq \Pr \left[\mathbf{Exp}_{\text{PKE},A}^{\text{ind-cca-BP}}(k) \Rightarrow \text{true} \wedge \text{BD} \right] + \mathbf{Adv}_{f,I}^{\text{ow}}.$$

Combining this with (2) and (3), we finally get (1). ■

Remark. We stress that our adversary A against $\overline{\text{PKE}}$'s IND-CCA-SP security does not query its decryption oracle after receiving the challenge ciphertext. Hence, $\overline{\text{PKE}}$ is not even IND-CCA1 secure. (Here IND-CCA1 security is defined like IND-CCA-SE security, except that the second stage A_2 of the adversary does not get access to a decryption oracle [31, 4].) Since any reasonable form of (full) IND-CCA security should imply IND-CCA1 security, we view this as another indication that IND-CCA-SE security is the “right” definition of IND-CCA security.

IND-CCA-SE \Rightarrow IND-CCA-SP. We already noted that IND-CCA-SP implies IND-CCA-SE. Theorem 3.2 below says that the converse is true as well, meaning that in the case where decryption of the challenge ciphertext is disallowed only in the second phase, the exclusion and penalty style notions are equivalent. (We will see below that this is not true in the case where the decryption of the challenge ciphertext is disallowed in both phases.) Theorem 3.2 is in fact understood in folklore but we state and prove it for completeness.

Theorem 3.2 [IND-CCA-SE \Rightarrow IND-CCA-SP] If PKE is IND-CCA-SE secure then PKE is IND-CCA-SP secure.

Proof: Given an adversary $A \in \mathcal{A}_{\text{PKE}}^{\text{SP}}$ against IND-CCA-SP security of PKE we show how to build an adversary $B \in \mathcal{A}_{\text{PKE}}^{\text{SE}}$ against IND-CCA-SE security of PKE such that for all $k \in \mathbb{N}$,

$$\mathbf{Adv}_{\text{PKE},A}^{\text{ind-cca-SP}}(k) \leq \mathbf{Adv}_{\text{PKE},B}^{\text{ind-cca-SE}}(k). \quad (8)$$

We let $B_1 = A_1$. Algorithm B_2 , given C^*, St , runs A_2 on C^*, St , and finally returns whatever A_2 returns. B_2 responds to A_2 's oracle queries as follows. When A_2 makes a query C , if $C \neq C^*$, B_2 responds with its own decryption oracle, else it returns \perp to A_2 . This ensures that in $\mathbf{Exp}_{\text{PKE},B}^{\text{ind-cca-SE}}$, we have $C^* \notin S_2$ with probability 1. Hence $B \in \mathcal{A}_{\text{PKE}}^{\text{SE}}$. Furthermore, Equation (8) holds since a decryption query satisfying $C = C^*$ directly implies that A loses. ■

IND-CCA-BE $\not\Rightarrow$ IND-CCA-BP. Our final separation shows that in the case where decryption of the challenge ciphertext is disallowed in both phases, the exclusion and penalty style notions are *not* equivalent. (This is in contrast to the case where decryption of the challenge ciphertext is disallowed only in the second phase, as noted above.)

Theorem 3.3 [IND-CCA-BE $\not\Rightarrow$ IND-CCA-BP] Assume there exist one-way permutations and a scheme PKE which is IND-CCA-BE secure. Then there exists a scheme $\overline{\text{PKE}}$ which is IND-CCA-BE secure but not secure in the sense of IND-CCA-BP.

Proof: Let $f : \{0,1\}^* \rightarrow \{0,1\}^*$ be a one-way permutation and assume that $\text{PKE} = (\text{Kg}, \text{Enc}, \text{Dec})$ is IND-CCA-BE secure. Consider the scheme $\overline{\text{PKE}} = (\overline{\text{Kg}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ of Figure 4 with $N_k = \{0,1\}^k$. First we show that $\overline{\text{PKE}}$ is not IND-CCA-BP secure. Adversary $A = (A_1, A_2)$ against $\overline{\text{PKE}}$ proceeds as follows. Given $\overline{pk} = (pk, Y)$, adversary A_1 queries $\text{DEC}_1(\cdot)$ on ciphertext $(1, 1^k)$ to obtain M_{weak} . It picks $M_1 \leftarrow_{\text{R}} \{0,1\}^k \setminus \{M_{\text{weak}}\}$ and returns $M_0 = M_{\text{weak}}$ and M_1 as the two challenge messages to the experiment. A_2 obtains a challenge ciphertext \overline{C}^* which is parsed as (s, C) . It returns $b' = 0$ if $s = 1$, and $b' = 1$ otherwise. Adversary A wins with probability 1 as long as $\overline{C}^* \notin S_1$ which happens with probability $1 - 2^{-k}$. Hence $\mathbf{Adv}_{\overline{\text{PKE}},A}^{\text{ind-cca-BP}}(k) = 1 - 2^{-k}$.

Note that the above adversary A is not contained in $\mathcal{A}_{\overline{\text{PKE}}}^{\text{BE}}$ since, with probability 2^{-k} , we have $\overline{C}^* \in S_1$. Indeed, we can show that $\overline{\text{PKE}}$ is IND-CCA-BE secure. The idea is again that an adversary needs to use M_{weak} as one of the challenge messages in order to win. However, an adversary from $\mathcal{A}_{\overline{\text{PKE}}}^{\text{BE}}$ using M_{weak} as one of the challenge messages can *never* make a decryption query \overline{C} of the form $(1, C)$ in the first phase, since $\overline{C}^* = (1, C)$ with non-zero probability $2^{-k}/2$. Hence, M_{weak} remains hidden through the one-way permutation. Details are similar to the proof of Claim 3 and omitted here. ■

<p>Experiment $\mathbf{Exp}_{\text{KEM},A}^{\text{ind-cca-X}}(k)$</p> <p>$(pk, sk) \leftarrow_{\text{R}} \text{Kg}(k); S_1, S_2 \leftarrow \emptyset$</p> <p>$St \leftarrow_{\text{R}} A_1^{\text{DEC}_1(\cdot)}(pk)$</p> <p>$(C^*, K_1^*) \leftarrow_{\text{R}} \text{Enc}(pk); K_0^* \leftarrow_{\text{R}} \mathcal{K}(k)$</p> <p>$b \leftarrow_{\text{R}} \{0, 1\}; b' \leftarrow_{\text{R}} A_2^{\text{DEC}_2(\cdot)}(C^*, K_b^*, St)$</p> <p>Return:</p> <p>SE, BE : $(b = b')$</p> <p>SP : $(b = b') \wedge (C^* \notin S_2)$</p> <p>BP : $(b = b') \wedge (C^* \notin S_1 \cup S_2)$</p>	<p>Oracle $\text{DEC}_1(C)$</p> <p>$S_1 \leftarrow S_1 \cup \{C\}$</p> <p>return $\text{Dec}(sk, C)$</p> <p>Oracle $\text{DEC}_2(C)$</p> <p>$S_2 \leftarrow S_2 \cup \{C\}$</p> <p>return $\text{Dec}(sk, C)$</p>
--	---

Figure 6: Experiment $\mathbf{Exp}_{\text{KEM},A}^{\text{ind-cca-X}}(k)$, for $X \in \{\text{SE}, \text{BE}, \text{SP}, \text{BP}\}$.

4 Results for Key Encapsulation Schemes

SYNTAX. A key space \mathcal{K} is a map that associates to any $k \in \mathbb{N}$ a finite set $\mathcal{K}(k) \subseteq \{0, 1\}^*$ of strings. The elements of $\mathcal{K}(k)$ are called keys, and it is required that $|\mathcal{K}(k)| \geq 2$ for all $k \in \mathbb{N}$. A key-encapsulation mechanism $\text{KEM} = (\text{Kg}, \text{Enc}, \text{Dec})$ over \mathcal{K} is a triple of algorithms. The key generation algorithm Kg takes a security parameter 1^k and returns a pair (pk, sk) of matching public and secret keys. The encapsulation algorithm Enc takes pk and produces a key $K \in \mathcal{K}(k)$ together with an encapsulated ciphertext C . The deterministic decapsulation algorithm Dec takes sk and C to produce either a key $K \in \mathcal{K}(k)$ or a special symbol \perp to indicate that the ciphertext was invalid. The consistency requirement is that for all $k \in \mathbb{N}$, for all (pk, sk) which can be output by $\text{Kg}(1^k)$ and for all (C, K) that can be output by $\text{Enc}(pk)$, we have that $\text{Dec}(sk, C) = K$.

IND-CCA SECURITY. A KEM IND-CCA adversary $A = (A_1, A_2)$ is a pair of algorithms. Let \mathcal{B} be the class of all such adversaries. Let $X \in \{\text{SP}, \text{BP}, \text{SE}, \text{BE}\}$. To an adversary $A = (A_1, A_2)$ and a KEM scheme KEM , we associate the experiment $\mathbf{Exp}_{\text{KEM},A}^{\text{ind-cca-X}}(k)$ in Figure 6. We define the advantage of A in the experiment as

$$\mathbf{Adv}_{\text{KEM},A}^{\text{ind-cca-X}}(k) = 2 \Pr[\mathbf{Exp}_{\text{KEM},A}^{\text{ind-cca-X}}(k) \Rightarrow \text{true}] - 1.$$

Let $\mathcal{B}_{\text{KEM}}^{\text{SP}} = \mathcal{B}_{\text{KEM}}^{\text{BP}} = \mathcal{B}$ be the class of all IND-CCA adversaries. Let $\mathcal{B}_{\text{KEM}}^{\text{SE}}$ be the class of all $A \in \mathcal{B}$ such that for all $k \in \mathbb{N}$, the probability that $C^* \in S_2$ in $\mathbf{Exp}_{\text{KEM},A}^{\text{ind-cca-SE}}(k)$ is 0. Let $\mathcal{B}_{\text{KEM}}^{\text{BE}}$ be the class of all $A \in \mathcal{B}$ such that for all $k \in \mathbb{N}$, the probability that $C^* \in S_1 \cup S_2$ in $\mathbf{Exp}_{\text{KEM},A}^{\text{ind-cca-BE}}(k)$ is 0. We say that KEM is IND-CCA- X secure if $\mathbf{Adv}_{\text{KEM},A}^{\text{ind-cca-X}}(\cdot)$ is negligible for all $A \in \mathcal{B}_{\text{KEM}}^X$.

We also consider the following simpler one-phase notions. A one-phase KEM IND-CCA adversary A consists of a single algorithm. Let $X \in \{\text{OP}, \text{OE}\}$. To an adversary A and KEM, we associate the one-phase experiment $\mathbf{Exp}_{\text{KEM},A}^{\text{ind-cca-X}}(k)$ in Figure 7. We define the advantage of A as above. Let $\mathcal{B}_{\text{KEM}}^{\text{OP}}$ be the class of all one-phase KEM IND-CCA adversaries. Let $\mathcal{B}_{\text{KEM}}^{\text{OE}}$ be the class of all $A \in \mathcal{B}_{\text{KEM}}^{\text{OP}}$ such that for all $k \in \mathbb{N}$, the probability that $C^* \in S$ in $\mathbf{Exp}_{\text{KEM},A}^{\text{ind-cca-OE}}(k)$ is 0. We say that KEM is IND-CCA- X secure if $\mathbf{Adv}_{\text{KEM},A}^{\text{ind-cca-X}}(\cdot)$ is a negligible function for all $A \in \mathcal{B}_{\text{KEM}}^X$.

SMOOTHNESS. For $k \in \mathbb{N}$ we let

$$\mathbf{Smth}_{\text{KEM}}(k) = \mathbf{E} \left[\max_{C \in \{0,1\}^*} \Pr_{(K,C') \leftarrow_{\text{R}} \text{Enc}(pk)} [C' = C] \right]$$

where the expected value is taken over all $(pk, sk) \leftarrow_{\text{R}} \text{Kg}(k)$. We refer to $\mathbf{Smth}_{\text{KEM}}(\cdot)$ as the smoothness of KEM and say that KEM is smooth if $\mathbf{Smth}_{\text{KEM}}(\cdot)$ is negligible. The notion of a smooth KEM scheme

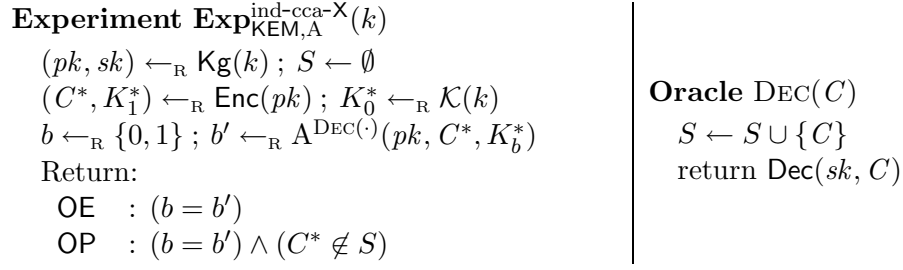


Figure 7: One-phase experiment $\text{Exp}_{\text{KEM},A}^{\text{ind-cca-X}}(k)$, for $X \in \{\text{OE}, \text{OP}\}$.




Figure 8: Relations between an expanded set of IND-CCA security notions for KEMs. The dotted lines are trivial implications, and the numbers annotating the solid line implications indicate the theorems establishing them.

will play a crucial role in the proof of Theorem 4.4 and may be of independent interest.²

RESULTS. Figure 8 depicts our results, which show that all six notions of IND-CCA-security for KEMs are equivalent. The equivalences of the right-hand-side of Figure 2 are a consequence. The trivial implications (dashed arrows) of Figure 8 should be clear from the definitions. We now prove the two other implications.

IND-CCA-OE \Rightarrow IND-CCA-BP. Theorem 4.1 below shows that security under the one-phase exclusion-style notion implies security under the two-phase penalty-style notion that disallows challenge-decryption in both phases. The (simple) proof of the following is given in Appendix B.

Theorem 4.1 [IND-CCA-OE \Rightarrow IND-CCA-BP] If KEM is IND-CCA-OE secure then KEM is IND-CCA-BP secure.

IND-CCA-BP \Rightarrow IND-CCA-SP. Theorem 4.4 below shows that for penalty-based notions allowing or disallowing a challenge-ciphertext query in the first phase does not make a difference. First, a useful lemma shows that for smooth KEMs, IND-CCA-BP security and IND-CCA-SP security are indeed equivalent.

Lemma 4.2 If KEM is smooth and IND-CCA-BP secure then it is IND-CCA-SP secure.

Proof: Given an adversary $A = (A_1, A_2) \in \mathcal{B}_{\text{KEM}}^{\text{SP}} = \mathcal{B}_{\text{KEM}}^{\text{BP}}$ we show that for all $k \in \mathbb{N}$,

$$\text{Adv}_{\text{KEM},A}^{\text{ind-cca-SP}}(k) \leq \text{Adv}_{\text{KEM},A}^{\text{ind-cca-BP}}(k) + Q_1(k) \cdot \text{Smth}_{\text{KEM}}(k), \quad (9)$$

where $Q_1(k)$ is a polynomial upper bound on the number of queries that A_1 makes. Details are similar to the proof of Theorem A.1 and omitted here. ■

²In fact, Fujisaki and Okamoto used essentially the same notion (called γ -uniformity in their work) in their result [19]; the main difference to our notion is the technicality that they quantify over all (pk, sk) , where we only consider the expected value over (pk, sk) .

Next we show that for KEM schemes $\overline{\text{IND-CCA-BP}}$ security implies smoothness. This is in contrast to PKE schemes where the counterexample $\overline{\text{PKE}}$ from Figure 4 shows a smooth PKE scheme which is not $\overline{\text{IND-CCA-BP}}$ secure.

Lemma 4.3 If KEM is $\overline{\text{IND-CCA-BP}}$ secure, then it is smooth.

Proof: We show that there exists an adversary $B = (B_1, B_2) \in \mathcal{B}_{\text{KEM}}^{\text{BP}}$ such that for all $k \in \mathbb{N}$,

$$\mathbf{Adv}_{\text{KEM},B}^{\text{ind-cca-BP}}(k) \geq \frac{1}{2} \cdot \mathbf{Smth}_{\text{KEM}}^2(k). \quad (10)$$

Adversary B_1 obtains $1^k, pk$ and returns $St = pk$. Adversary B_2 obtains (pk, C^*, K^*) and proceeds as follows. It picks random $(K', C') \leftarrow_{\text{R}} \text{Enc}(pk)$. If $C^* \neq C'$ then B_2 picks a random bit b' and returns it. If $C^* = C'$ then B_2 returns $b' = 1$ if $K' = K^*$ and $b' = 0$, otherwise.

We now turn to the analysis of B . For any pk and $C \in \{0, 1\}^*$ let

$$\nu(pk, C) = \Pr_{(\tilde{K}, \tilde{C}) \leftarrow_{\text{R}} \text{Enc}(pk)} [\tilde{C} = C]$$

Let $C_{\max}(pk)$ be such that $\nu(pk, C_{\max}(pk)) \geq \nu(pk, C)$ for all $C \in \{0, 1\}^*$. We define GD as the event that $C' = C_{\max}(pk)$ and $C^* = C_{\max}(pk)$ in $\mathbf{Exp}_{\text{KEM},B}^{\text{ind-cca-BP}}(k)$. Assume GD has happened and hence $C^* = C'$. If $b = 1$ then B wins with probability 1 since (by consistency) $K^* = K'$. If $b = 0$ then B only loses if the two keys K' and K^* collide. Since the experiment picks $K^* = K_0^*$ uniformly distributed from $\mathcal{K}(k)$ this happens with probability $1/|\mathcal{K}(k)| \leq 1/2$.

$$\Pr[b = b' \mid \text{GD}] = \frac{1}{2} \cdot (\Pr[b = b' \mid \text{GD} \wedge b = 0] + \Pr[b = b' \mid \text{GD} \wedge b = 1]) \geq \frac{1}{2}(1 + 1 - \frac{1}{2}) = \frac{3}{4}.$$

On the other hand, $\Pr[b = b' \mid \neg \text{GD}] \geq 1/2$. Since B never queries the decapsulation oracle we have

$$\begin{aligned} \mathbf{Adv}_{\text{KEM},B}^{\text{ind-cca-BP}}(k) &= 2 \Pr[\mathbf{Exp}_{\text{KEM},B}^{\text{ind-cca-BP}}(k) \Rightarrow \mathbf{true}] - 1 = 2 \Pr[b = b'] - 1 \\ &\geq 2(\Pr[b = b' \mid \text{GD}] \cdot \Pr[\text{GD}] + \Pr[b = b' \mid \neg \text{GD}] \cdot (1 - \Pr[\text{GD}])) - 1 \\ &\geq \frac{1}{2} \cdot \Pr[\text{GD}] \end{aligned}$$

It remains to bound $\Pr[\text{GD}]$. To this end let

$$X(pk) = \Pr_{(K,C) \leftarrow_{\text{R}} \text{Enc}(pk)} [C = C_{\max}(pk)].$$

Regard X as a random variable over the choice of pk given by $(pk, sk) \leftarrow_{\text{R}} \text{Kg}(1^k)$. Then, taking the expectation over the choice of (pk, sk) we have $\mathbf{E}[X] \geq \mathbf{Smth}_{\text{PKE}}(k)$ so

$$\Pr[\text{GD}] = \mathbf{E}[X^2] \geq \mathbf{E}[X]^2 \geq \mathbf{Smth}_{\text{PKE}}^2(k)$$

due to Jensen's inequality. This yields Equation (10) and concludes the proof of the claim. \blacksquare

The preceding two lemmas can be combined to show our main result for KEMs:

Theorem 4.4 [$\overline{\text{IND-CCA-BP}} \Rightarrow \overline{\text{IND-CCA-SP}}$] If KEM is $\overline{\text{IND-CCA-BP}}$ secure then KEM is $\overline{\text{IND-CCA-SP}}$ secure.

Proof: Combining Lemma 4.2 and Lemma 4.3 (and in particular, Equations (9) and (10)), we obtain that for all $A \in \mathcal{A}_{\text{KEM}}^{\text{SP}}$ there exists $B \in \mathcal{B}_{\text{KEM}}^{\text{BP}}$ such that for all $k \in \mathbb{N}$,

$$\mathbf{Adv}_{\text{KEM},B}^{\text{ind-cca-BP}}(k) \geq \frac{1}{2} \cdot \left(\frac{\mathbf{Adv}_{\text{KEM},A}^{\text{ind-cca-SP}}(k)}{Q_1(k) + 1} \right)^2. \quad (11)$$

This proves the theorem. \blacksquare

References

- [1] Michel Abdalla, Dario Catalano, Alex Dent, John Malone-Lee, Gregory Neven, and Nigel Smart. Identity-based encryption gone wild. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006, Part II*, volume 4052 of *LNCS*, pages 300–311. Springer-Verlag, Berlin, Germany, July 2006. (Cited on page 3.)
- [2] Masayuki Abe. Combining encryption and proof of knowledge in the random oracle model. *Comput. J.*, 47(1):58–70, 2004. (Cited on page 3.)
- [3] Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 259–274. Springer-Verlag, Berlin, Germany, May 2000. (Cited on page 3.)
- [4] Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. Relations among notions of security for public-key encryption schemes. In Hugo Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 26–45. Springer-Verlag, Berlin, Germany, August 1998. (Cited on page 3, 4, 9.)
- [5] Daniel Bleichenbacher. Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. In Hugo Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 1–12. Springer-Verlag, Berlin, Germany, August 1998. (Cited on page .)
- [6] Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM Journal on Computing*, 36(5):1301–1328, 2007. (Cited on page 4.)
- [7] Dan Boneh and Matthew K. Franklin. Identity-based encryption from the Weil pairing. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 213–229. Springer-Verlag, Berlin, Germany, August 2001. (Cited on page 4.)
- [8] Xavier Boyen, Qixiang Mei, and Brent Waters. Direct chosen ciphertext security from identity-based techniques. In Vijayalakshmi Atluri, Catherine Meadows, and Ari Juels, editors, *ACM CCS 05*, pages 320–329. ACM Press, November 2005. (Cited on page 3.)
- [9] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001. (Cited on page 4.)
- [10] Ran Canetti, Hugo Krawczyk, and Jesper Buus Nielsen. Relaxing chosen-ciphertext security. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 565–582. Springer-Verlag, Berlin, Germany, August 2003. (Cited on page 4.)
- [11] Benoît Chevallier-Mames, Duong Hieu Phan, and David Pointcheval. Optimal asymmetric encryption and signature paddings. In John Ioannidis, Angelos Keromytis, and Moti Yung, editors, *ACNS 05*, volume 3531 of *LNCS*, pages 254–268. Springer-Verlag, Berlin, Germany, June 2005. (Cited on page 3.)
- [12] Jean-Sébastien Coron, Helena Handschuh, Marc Joye, Pascal Paillier, David Pointcheval, and Christophe Tymen. Optimal chosen-ciphertext secure encryption of arbitrary-length messages. In David Naccache and Pascal Paillier, editors, *PKC 2002*, volume 2274 of *LNCS*, pages 17–33. Springer-Verlag, Berlin, Germany, February 2002. (Cited on page 3.)
- [13] Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Hugo Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 13–25. Springer-Verlag, Berlin, Germany, August 1998. (Cited on page 3, 4, 16.)

- [14] Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM Journal on Computing*, 33(1):167–226, 2003. (Cited on page 3.)
- [15] Hans Delfs and Helmut Knebel. *Introduction to Cryptography*. Springer Verlag, second edition, 2007. (Cited on page 3.)
- [16] Alex Dent. A designer’s guide to KEMs. In K. G. Paterson, editor, *Cryptography and Coding, 9th IMA International Conference*, volume 2898 of *LNCS*, pages 133–151, Cirencester, UK, 2003. Springer-Verlag, Berlin, Germany. (Cited on page 3.)
- [17] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000. (Cited on page 4.)
- [18] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In G. R. Blakley and David Chaum, editors, *CRYPTO’84*, volume 196 of *LNCS*, pages 10–18. Springer-Verlag, Berlin, Germany, August 1985. (Cited on page 16.)
- [19] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Michael J. Wiener, editor, *CRYPTO’99*, volume 1666 of *LNCS*, pages 537–554. Springer-Verlag, Berlin, Germany, August 1999. (Cited on page 3, 11.)
- [20] Eiichiro Fujisaki, Tatsuaki Okamoto, David Pointcheval, and Jacques Stern. RSA-OAEP is secure under the RSA assumption. *Journal of Cryptology*, 17(2):81–104, March 2004. (Cited on page 3.)
- [21] Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004. (Cited on page 3.)
- [22] Jens Groth. Rerandomizable and replayable adaptive chosen ciphertext attack secure cryptosystems. In Moni Naor, editor, *Theory of Cryptography, Proceedings of TCC 2004*, number 2951 in *Lecture Notes in Computer Science*, pages 152–170. Springer-Verlag, 2004. (Cited on page 4.)
- [23] Dennis Hofheinz, Jörn Müller-Quade, and Rainer Steinwandt. On modeling IND-CCA security in cryptographic protocols. *Tatra Mountains Mathematical Publications*, 33:83–97, 2006. (Cited on page 4.)
- [24] Jonathan Katz and Juhuda Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC Press, aug 2007. (Cited on page 3.)
- [25] Eike Kiltz. Chosen-ciphertext security from tag-based encryption. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 581–600. Springer-Verlag, Berlin, Germany, March 2006. (Cited on page 3.)
- [26] Eike Kiltz. Chosen-ciphertext secure key-encapsulation based on Gap Hashed Diffie-Hellman. In *Proceedings of PKC 2007*, volume 4450 of *LNCS*, pages 282 – 297, 2007. Full version available from <http://eprint.iacr.org/2007/036>. (Cited on page 3.)
- [27] Eike Kiltz and David Galindo. Direct chosen-ciphertext secure identity-based key encapsulation without random oracles. In *ACISP 2006*, volume 4058 of *LNCS*. Springer-Verlag, 2006. (Cited on page 4.)
- [28] Kaoru Kurosawa and Yvo Desmedt. A new paradigm of hybrid encryption scheme. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 426–442. Springer-Verlag, Berlin, Germany, August 2004. (Cited on page 3.)

- [29] Yehuda Lindell. A simpler construction of CCA2-secure public-key encryption under general assumptions. *Journal of Cryptology*, 19(3):359–377, July 2006. (Cited on page 3.)
- [30] Alfred J. Menezes, Paul C. van Oorschot, and Scott A. Vanstone. *Handbook of Applied Cryptography*. The CRC Press series on discrete mathematics and its applications. CRC Press, 2000 N.W. Corporate Blvd., Boca Raton, FL 33431-9868, USA, 1997. (Cited on page 3.)
- [31] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd ACM STOC*. ACM Press, May 1990. (Cited on page 4, 9.)
- [32] Tatsuaki Okamoto and David Pointcheval. REACT: Rapid Enhanced-security Asymmetric Cryptosystem Transform. In David Naccache, editor, *CT-RSA 2001*, volume 2020 of *LNCS*, pages 159–175. Springer-Verlag, Berlin, Germany, April 2001. (Cited on page 3.)
- [33] Pascal Paillier and Jorge L. Villar. Trading one-wayness against chosen-ciphertext security in factoring-based encryption. In Xuejia Lai and Kefei Chen, editors, *ASIACRYPT 2006*, volume 4284 of *LNCS*, pages 252–266. Springer-Verlag, Berlin, Germany, December 2006. (Cited on page 3.)
- [34] Duong Hieu Phan and David Pointcheval. On the security notions for public-key encryption schemes. In Carlo Blundo and Stelvio Cimato, editors, *SCN 04*, volume 3352 of *LNCS*, pages 33–46. Springer-Verlag, Berlin, Germany, September 2004. (Cited on page 3, 4.)
- [35] Manoj Prabhakaran and Mike Rosulek. Rerandomizable RCCA encryption. In Alfred Menezes, editor, *CRYPTO 2007*, *LNCS*, pages 517–534. Springer-Verlag, Berlin, Germany, August 2007. (Cited on page 4.)
- [36] Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In Joan Feigenbaum, editor, *CRYPTO’91*, volume 576 of *LNCS*, pages 433–444. Springer-Verlag, Berlin, Germany, August 1992. (Cited on page 4.)
- [37] Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th FOCS*, pages 543–553. IEEE Computer Society Press, October 1999. (Cited on page 3.)
- [38] Victor Shoup. Why chosen ciphertext security matters. IBM Research Report RZ 3076, November 1998. (Cited on page .)
- [39] Victor Shoup. OAEP reconsidered. *Journal of Cryptology*, 15(4):223–249, 2002. (Cited on page 3.)
- [40] Victor Shoup. ISO 18033-2: An emerging standard for public-key encryption. <http://shoup.net/iso/std6.pdf>, December 2004. Final Committee Draft. (Cited on page 3.)
- [41] Nigel P. Smart. The exact security of ECIES in the generic group model. In B. Honary, editor, *Cryptography and Coding, 8th IMA International Conference*, volume 2260 of *LNCS*, pages 73–84. Springer-Verlag, Berlin, Germany, 2001. (Cited on page 3.)

A Relations for smooth PKE schemes

We mentioned earlier some intuition for why one might think that disallowing decryption of the challenge ciphertext in both phases is equivalent to disallowing it only in the second phase, namely that, even for IND-CPA schemes, there must be, for every message, a large number of corresponding ciphertexts, and hence an adversary would be unable to predict (and hence query) the challenge ciphertext in the first phase. The counter-example of Theorem 3.1 shows this intuition is false in general; in the scheme $\overline{\text{PKE}}$ we built there, there is a message, namely M_{weak} , encryption of which can result in just one ciphertext, and yet the scheme is IND-CCA-BP (and hence IND-CPA) secure but not IND-CCA-SP




Figure 9: Implications and separations between the various IND-CCA security notions for PKE schemes with smooth ciphertexts.

secure. However, we now claim that the basic intuition mentioned above is still right in the sense that if indeed, for every message, there is a large number of corresponding ciphertexts —we will call this property smoothness— then indeed IND-CCA-BP implies IND-CCA-SP. Where the intuition went wrong was in thinking smoothness is implied by security properties like IND-CPA or IND-CCA-BP. (The scheme of Theorem 3.1 shows it is not.) Interestingly, we will however see that IND-CCA-BE and IND-CCA-SE are not equivalent even for smooth schemes, indicating the weakness of exclusion-based definitions. To detail all this we now define smoothness formally. For any $k \in \mathbb{N}$ and any scheme $\text{PKE} = (\text{Kg}, \text{Enc}, \text{Dec})$, we let

$$\mathbf{Smath}_{\text{PKE}}(k) = \mathbf{E} \left[\max_{M \in \{0,1\}^*, C \in \{0,1\}^*} \Pr_{C' \leftarrow_{\text{R}} \text{Enc}(pk, M)} [C = C'] \right]$$

where the expected value is taken over all $(pk, sk) \leftarrow_{\text{R}} \text{Kg}(k)$. We refer to $\mathbf{Smath}_{\text{PKE}}(k)$ as the smoothness of PKE and say that PKE is *smooth* if $\mathbf{Smath}_{\text{PKE}}(\cdot)$ is negligible.

Smooth practical schemes include the ElGamal scheme [18] and the Cramer-Shoup scheme [13]. For these schemes, $\mathbf{Smath}_{\text{PKE}}(k) \leq 2^{-k}$. On the other hand, the scheme $\overline{\text{PKE}}$ from Theorem 3.1 is not smooth: For any (pk, sk) , for the message M_{weak} and the ciphertext $C = (1, 1)$ we have $\Pr[C = \text{Enc}(pk, M_{\text{weak}})] = 1$ so $\mathbf{Smath}_{\overline{\text{PKE}}}(k) = 1$. The relations between the different IND-CCA notions for PKE schemes with smooth ciphertexts are summarized in Figure 9. The difference between this and Figure 2 is that IND-CCA-BP now implies IND-CCA-SP.

Theorem A.1 If the scheme PKE is IND-CCA-BP secure and smooth, then it is also IND-CCA-SP secure.

Proof: Given an adversary $A = (A_1, A_2) \in \mathcal{A}_{\text{PKE}}^{\text{SP}} = \mathcal{A}_{\text{PKE}}^{\text{BP}}$ we show that for all $k \in \mathbb{N}$,

$$\mathbf{Adv}_{\text{PKE}, A}^{\text{ind-cca-SP}}(k) \leq \mathbf{Adv}_{\text{PKE}, A}^{\text{ind-cca-BP}}(k) + 2Q_1(k) \cdot \mathbf{Smath}_{\text{PKE}}(k), \quad (12)$$

where $Q_1(k)$ is a polynomial upper bound on the number of decryption queries of A_1 .

We define the event BD in $\mathbf{Exp}_{\text{PKE}, A}^{\text{ind-cca-BP}}$ to hold when $C^* \in S_1$. Then

$$\mathbf{Adv}_{\text{PKE}, A}^{\text{ind-cca-SP}}(k) \leq \mathbf{Adv}_{\text{PKE}, A}^{\text{ind-cca-BP}}(k) + 2 \Pr[\text{BD}]. \quad (13)$$

On the other hand we have $\Pr[\text{BD}] \leq Q_1(k) \cdot \mathbf{Smath}_{\text{PKE}}(k)$ because for any given first phase query C , the smoothness property of PKE guarantees that $\Pr[C = C^*] \leq \mathbf{Smath}_{\text{PKE}}(k)$. Finally, a union bound leads to the claimed statement. ■

However, Theorem A.2 below shows that, even for smooth schemes, the equivalence between allowing challenge decryption queries in both or just the second phase does *not* carry over to the case of exclusion-based definitions.

Theorem A.2 [IND-CCA-BE $\not\equiv$ IND-CCA-SE] Assume there exist one-way permutations and a smooth scheme PKE which is IND-CCA-BE secure. Then there exists a smooth scheme $\overline{\text{PKE}}$ which is IND-CCA-BE secure but not IND-CCA-SE secure.

Proof: Assume PKE is IND-CCA-BE secure and smooth. We use the IND-CCA-BE secure PKE scheme $\overline{\text{PKE}}$ from the the proof of Theorem 3.3 (Figure 4 with $N_k = \{0,1\}^k$). Note that $\mathbf{Smth}_{\overline{\text{PKE}}}(k) \leq \mathbf{Smth}_{\text{PKE}}(k) + 2^{-k}$ and hence $\overline{\text{PKE}}$ is smooth.

Consider the adversary $A = (A_1, A_2)$ used in the proof of Theorem 3.3 to attack IND-CCA-BP security of the scheme. Since A_2 never queries the decryption oracle we have that $A \in \mathcal{A}_{\overline{\text{PKE}}}^{\text{SE}}$. Furthermore, A wins with probability 1, always, and hence $\overline{\text{PKE}}$ is not IND-CCA-SE secure. ■

B Proof of Theorem 4.1

Proof: Let $B = (B_1, B_2) \in \mathcal{B}_{\text{KEM}}^{\text{BP}}$. We build an adversary $A \in \mathcal{B}_{\text{KEM},A}^{\text{OE}}$ such that for all $k \in \mathbb{N}$,

$$\mathbf{Adv}_{\text{KEM},B}^{\text{ind-cca-BP}}(k) \leq \mathbf{Adv}_{\text{KEM},A}^{\text{ind-cca-OE}}(k). \quad (14)$$

A obtains $(1^k, pk, C^*, K_b^*)$ and runs B_1 on $(1^k, pk)$ and inputs St . Next, A runs B_2 on input (St, C^*, K_b^*) and outputs whatever B_2 returns. During the executions, A needs to answer B_1 and B_2 's decapsulation queries. Let C be such a decapsulation query made by B_1 or B_2 . If $C \neq C^*$ then A answers using its own decapsulation oracle. If $C = C^*$ is queried, then A aborts. This implies Equation (14) since a successful adversary $B \in \mathcal{B}_{\text{KEM}}^{\text{BP}}$ is obliged not to submit C^* to the decapsulation oracle *at any time*. Furthermore, by construction, $A \in \mathcal{B}_{\text{KEM}}^{\text{OE}}$ which proves the theorem. ■