Parallel Decryption Queries in Bounded Chosen Ciphertext Attacks

Takahiro Matsuda* and Kanta Matsuura

The University of Tokyo, Japan
{tmatsuda,kanta}@iis.u-tokyo.ac.jp

Abstract. Whether it is possible to construct a chosen ciphertext secure (CCA secure) public key encryption (PKE) scheme only from a chosen plaintext secure (CPA secure) one is a fundamental open problem, and the best known positive results regarding this problem are the constructions of so-called *bounded CCA* secure schemes. Since we can achieve the best possible security in the bounded CCA security notions, in order to further tackle the problem, we would need other new security notions that capture intermediate security notions that lie between CPA and CCA security. Motivated by this situation, we focus on “parallel” decryption queries (originally introduced by Bellare and Sahai) for the extension of bounded CCA security, and introduce a new security notion which we call *mixed CCA* security. It captures security against adversaries that make single and parallel decryption queries in a predetermined order, where each parallel query can contain *unboundedly* many ciphertexts. Moreover, how the decryption oracle is available before and after the challenge is also taken into account in this new security definition, which enables us to capture existing major security notions that lie between CPA and CCA security in a unified security notion. We investigate the relations among mixed CCA security notions, and show a necessary and sufficient condition of implications/separations between any two notions in mixed CCA security. We also show two black-box constructions of PKE schemes with improved security only using CPA secure schemes as building blocks.

Keywords: public key encryption, bounded CCA security, parallel decryption query, relations among security notions, black-box construction.

1 Introduction

Background. Studies on constructing and understanding public key encryption (PKE) schemes that satisfy security against chosen ciphertext attacks (CCA) [22,27], which is nowadays considered as a standard security notion needed in most practical applications/situations where PKE schemes are used, are important research topics in the area of cryptography. We can roughly categorize the

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

approaches for constructing CCA secure PKE schemes into two types: Constructions from specific number-theoretic assumptions and constructions from general assumptions. (In the following, we write IND-CCA1 to denote non-adaptive CCA security [22] and IND-CCA2 to denote adaptive CCA security [27]).

The approaches of the first type have been successful so far from both theoretical and practical points of view. After the first novel practical scheme based on the decisional Diffie-Hellman (DDH) assumption by Cramer and Shoup [10], many practical IND-CCA2 secure PKE schemes that pursue smaller ciphertext size and/or base security on weaker assumptions have been constructed so far, e.g. [20,6,17,7,14,18,15,30].

The approaches of the second type have also been successful, mainly from a theoretical point of view. Especially, it is known that if there exists an (enhanced) trapdoor permutation, which is one of the most fundamental primitives, then we can construct an IND-CCA2 secure PKE schemes generically [22,4,11]. There are also several elegant generic constructions of IND-CCA2 secure PKE schemes from primitives with some “stronger” functionality and/or security, such as constructions from identity-based encryption [5], and from special types of injective trapdoor functions and trapdoor relations [25,29,19,30].

However, one of the most fundamental problems still remains open: *Is it possible to generically construct a CCA (IND-CCA1 or IND-CCA2) secure PKE scheme from any semantically secure [13] (i.e. IND-CPA secure) one?*

So far, there are several negative and positive results related to this problem. Gertner et al. [12] showed that constructing an IND-CCA1 secure PKE scheme only from IND-CPA secure PKE schemes in a black-box manner is impossible, if the construction satisfies the property called *shielding*, where a PKE-to-PKE construction is said to be shielding if the decryption algorithm of the construction does not call the encryption algorithm of the building block PKE scheme.

Pass et al. [23] showed how to construct a PKE scheme that is non-malleable against chosen plaintext attacks (NM-CPA) from any IND-CPA secure PKE scheme. Their construction uses a certain class of NIZK proofs and was non-black-box.

Cramer et al. [9] introduced the notion of *bounded CCA* security which is defined in exactly the same way as ordinary IND-CCA2 security, except that the number of decryption oracle queries that an adversary can ask is bounded by some predetermined value (say, q) that is known a priori (we denote this notion by q -CCA2). Then they showed that for any polynomial q it is possible to construct an IND- q -CCA2 secure PKE scheme from any IND-CPA secure one in a black-box and shielding manner. They furthermore showed that for any polynomial q it is possible to construct a PKE scheme that satisfies non-malleability against q -bounded CCA (NM- q -CCA2) in a non-black-box manner.

Recently, Choi et al. [8] showed the constructions of PKE schemes from any IND-CPA secure scheme both in a black-box and shielding manner. Their first construction achieves NM-CPA security, and their second construction, which is essentially the same as the first construction but needs larger parameters, can achieve NM- q -CCA2 security.

These previous results show that we can achieve the best possible security notion (NM- q -CCA2) in the bounded CCA framework. This suggests that in order to proceed from the current situation, we would need new security notions which are intermediate between CPA and CCA security in a different sense from bounded CCA security. The motivation of this paper is to introduce and study such intermediate security notions as an extension of the bounded CCA security as a foundation for tackling the above fundamental problem.

Extending Bounded CCA Security with Parallel Decryption Queries. For the purpose mentioned above, we focus on and use the concept of the *parallel chosen ciphertext attacks* which is originally introduced by Bellare and Sahai [3] in the context of non-malleability [11] for PKE schemes, and consider *parallel* queries in the bounded CCA security framework. More specifically, as an extension of bounded CCA security, we introduce a new security notion, which we call *mixed CCA security*, that captures security against adversaries that make single (i.e. ordinary) decryption queries and parallel decryption queries in a predetermined order, where each parallel query can contain *unboundedly* many ciphertexts. (The name “mixed” is because we consider a mix of single and parallel queries). Moreover, the difference among decryption queries that are only allowed to make before/after the challenge and those that are allowed to make both before and after the challenge (an adversary can decide “flexibly” how to issue queries as long as it follows the predetermined order of queries and types) is also taken into account in our definition, which enables us to capture existing major security notions that lie between CPA and CCA security, including slightly complex notions such as non-malleability against bounded CCA (NM- q -CCA2) that considers “stage-specific” decryption queries, in a unified security notion. As a natural and interesting special class of mixed CCA security, we also introduce the notion of *bounded parallel CCA security*. For more details, see Section 3. We believe that the mixed CCA security provides a theoretical foundation for discussion of the problem of whether constructing (unbounded) CCA secure PKE schemes from any CPA secure PKE schemes is possible or not, and for intermediate results towards the problem.

1.1 Our Contributions

Relations among Mixed CCA Security Notions. We investigate the relations among mixed CCA security notions for PKE schemes and for key encapsulation mechanisms (KEMs) in Section 4. As one of the main results, we show necessary and sufficient conditions for implications/separations between any two notions in mixed CCA security. Interestingly and perhaps somewhat surprisingly, *the relations for PKE schemes differ depending on its plaintext space size*. More specifically, the relations among security notions for PKE schemes with super-polynomially large plaintext space size and those with polynomially bounded plaintext space size are different. Therefore, this difference suggests that when we consider the relations among security notions for PKE schemes, we have to be also careful about the plaintext space size, though seemingly unrelated.

The relations for KEMs are the same as those of PKE schemes with polynomially bounded plaintext space size.

Black-Box Feasibility Results from CPA-Security. Using the notion of mixed CCA security, in Section 5, we show two new black-box constructions of PKE schemes (which can encrypt plaintexts of polynomial length, and thus have exponentially large plaintext space size) from an IND-CPA secure PKE scheme. The first one is constructed based on the construction by Choi et al. [8] which is NM- q -CCA2 secure, and achieves slightly but strictly stronger security notion than NM- q -CCA2. Our approach for the first construction is to use the Choi et al. scheme as a KEM and combine it with an IND-CCA2 secure data encapsulation mechanism (DEM), and thus is a very simple extension. In order for this simple approach to work, we show some implication result for mixed CCA security of KEMs (and PKE schemes with polynomially bounded plaintext space size). The second one is constructed based on the above result and the construction of PKE scheme by Cramer et al. [9], and achieves yet another security notion which cannot be directly compared with the security notion achieved by our other constructions and with NM- q -CCA2 security.

As will be explained later, one of the important and interesting observations that our results suggests, combined with previously known results, is that *the difficulty of constructing an IND-CCA1 secure PKE scheme only from IND-CPA secure one lies not in whether the number of decryption results that an adversary can see is bounded or not, but in whether the number of the adversary's "adaptive" decryption queries is bounded.* To the best of our knowledge, this observation has not been explicitly stated before.

2 Preliminaries

In this section, we review the basic notations and definitions of primitives used in this paper. Due to space limitation, the definitions for key encapsulation mechanism (KEM) and data encapsulation mechanism (DEM) are omitted and are provided in the full version. (They can also be found in [10,16], for example).

Basic Notations. If q is a natural number, then $[q] = \{1, \dots, q\}$. " $x \leftarrow y$ " 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. " $|x|$ " denotes the size of the set if x is a finite set or bit length of x if x is an element of some set. "PPTA" denotes a *probabilistic polynomial time algorithm*. $\mathcal{A}^{\mathcal{O}}$ denotes an algorithm $\mathcal{A}$ with oracle access to $\mathcal{O}$. Unless otherwise stated, k denotes the security parameter. A function $f : \mathbb{N} \rightarrow [0, 1]$ is said to be *negligible* if for any positive polynomial $p(k)$ and for all sufficiently large k , we have $f(k) < \frac{1}{p(k)}$.

Public Key Encryption. A PKE scheme Π consists of the following three PPTAs (PKG, PEnc, PDec): A key generation algorithm PKG takes 1^k (security parameter k) as input, and outputs a public/private key pair (pk, sk) ; An encryption algorithm PEnc takes pk and a plaintext $m \in \mathcal{M}_{\Pi}$ (the plaintext space

of Π) as input, and outputs a ciphertext c ; A deterministic decryption algorithm PDec takes sk and c as input, and outputs a plaintext m (or a symbol $\perp$ meaning “decryption error”). As a correctness requirement, we require $\text{PDec}(sk, \text{PEnc}(pk, m)) = m$ for all (pk, sk) output from PKG and all $m \in \mathcal{M}_\Pi$.

Conventional Security Notions. The security notions for PKE schemes are expressed by a combination of a **GOAL** and an attack type (**ATK**) of an adversary. For conventional security notions for PKE schemes, we consider *indistinguishability* (**IND**) and *non-malleability* (**NM**) for security goals and *chosen plaintext attacks* (**CPA**), *non-adaptive chosen ciphertext attacks* (**CCA1**), *adaptive chosen ciphertext attacks* (**CCA2**), and *q-bounded chosen ciphertext attacks* (**q-CCA2**) [9] for attack types of an adversary. Non-malleability for PKE schemes we treat in this paper is the so-called *parallel chosen-ciphertext attack* based definition [3], which is equivalent to the indistinguishability based definition used in [23,24]¹. Since these conventional **GOAL-ATK** security notions can be expressed as special cases of *mixed CCA* security defined in Section 3, here we omit the definitions.

Implications and Separations of Security Notions. We will show several implications and separations of security notions, and thus we recall here. Though we write only the definition for PKE schemes, the same is defined for KEMs.

Definition 1. Let $\mathbf{X}$ and $\mathbf{Y}$ be security notions for PKE schemes. We say that $\mathbf{X}$ security implies $\mathbf{Y}$ security if any $\mathbf{X}$ secure PKE scheme is also $\mathbf{Y}$ secure. We say that $\mathbf{X}$ security does not imply $\mathbf{Y}$ security if, under the assumption that $\mathbf{X}$ secure PKE schemes exist, there exists a PKE scheme which is $\mathbf{X}$ secure but is not $\mathbf{Y}$ secure. We say that $\mathbf{X}$ security and $\mathbf{Y}$ security are equivalent if we have implications for both directions (i.e. from $\mathbf{X}$ to $\mathbf{Y}$ and from $\mathbf{Y}$ to $\mathbf{X}$).

Shielding Black-Box Constructions. We briefly recall the definition of a shielding black-box construction of an $\mathbf{X}$ secure PKE scheme from a $\mathbf{Y}$ secure scheme. The notion of black-box constructions we mention in this paper is classified as *fully-black-box* ones [28], but specified for PKE-to-PKE constructions. (for details, see [28]). The notion of the *shielding* constructions is from [12].

Definition 2. Let $\mathbf{X}$ and $\mathbf{Y}$ be security notions for PKE schemes. We say that there exists a shielding black-box construction of an $\mathbf{X}$ secure PKE scheme from a $\mathbf{Y}$ secure one, if there exist oracle PPTAs $\Pi = (\text{PKG}, \text{PEnc}, \text{PDec})$ and $\mathcal{B}$ with the following properties. For all algorithms $\pi = (G, E, D)$ and $\mathcal{A}$ (each algorithm can be of arbitrary complexity), the following two conditions are satisfied:

(**Correctness:**) If $\pi = (G, E, D)$ satisfies correctness as a PKE scheme, so does $\Pi^{G,E,D} = (\text{PKG}^{G,E,D}, \text{PEnc}^{G,E,D}, \text{PDec}^{G,D})$.

(**Security:**) If $\mathcal{A}$ breaks $\mathbf{X}$ security of $\Pi^{G,E,D} = (\text{PKG}^{G,E,D}, \text{PEnc}^{G,E,D}, \text{PDec}^{G,D})$ then $\mathcal{B}^{\mathcal{A},G,E,D}$ breaks $\mathbf{Y}$ security of π .

(Note that PDec does not have access to E).

¹ Pass et al. [24] prove that *many-message* (indistinguishability-based) non-malleability, which considers multiple challenge messages, and single-message non-malleability, considered in this paper, are equivalent.

3 Extending Bounded CCA: Mixed CCA Security

In order to deal with and discuss existing security notions for PKE schemes and KEMs that lie between IND-CPA and IND-CCA2 security in a unified way, in this section we introduce an extension of conventional bounded CCA security [9], which we call security against *mixed chosen ciphertext attacks* (mixed CCA security), where the decryption oracle in the security experiment accepts both single decryption queries and *parallel* decryption queries in a predetermined order, and “how” the decryption oracle is available before/after the challenge is also taken into account.

Preliminary Definitions. We first formally define the notion of a parallel query to an oracle.

Definition 3. Let $\mathcal{O} : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be an oracle. A parallel query to $\mathcal{O}$ is a vector $\vec{x} = (x_1, x_2, \dots)$ of inputs for $\mathcal{O}$, where the size of the vector $\vec{x}$ is not predetermined, and a response to the parallel query $\vec{x}$ is a vector of the output values $\vec{y} = (y_1, y_2, \dots)$ where $y_i = \mathcal{O}(x_i)$ for every $1 \leq i \leq |\vec{x}|$.

We stress that the number of inputs in each parallel query $\vec{x}$ is unbounded and can be dependent only on an algorithm that uses the oracle.

To define mixed CCA security, we need to introduce several notations. The symbols “s” and “p” denote one *single query* and one *parallel query*, respectively. Let $q \geq 0$ be an integer. “s^q” and “p^q” denote q single queries and q parallel queries, respectively. We define $s^0 = p^0 = \emptyset$.

If we write “(s^{q₁}p^{q₂}...)” with some integers $q_1, q_2, \dots \geq 0$, then it denotes a *query sequence*. This query sequence will define how the decryption oracle in the mixed CCA experiment accepts the queries. For example, (s²p³) denotes two single decryption queries followed by three parallel decryption queries. We denote by “unbound” a special sequence that indicates “unboundedly” many single queries, i.e. unbound = s[∞].

“QS” denotes a set consisting of all possible query sequences with the restriction that the total number of queries in each sequence is bounded to be polynomial (in the security parameter). We furthermore define $QS^* = QS \cup \{\text{unbound}\}$. We refer to queries following the query sequence $\text{seq} \in QS^*$ as “seq-queries”.

If $\text{seq} \in QS$, then we denote by “|seq|” the length of the query sequence. For example, if $\text{seq} = (s^2p)$ then $|\text{seq}| = 3$. We define $|\text{unbound}| = \infty$.

We define a concatenation operation “||” for query sequences naturally. For example, if $\text{seq}_1 = (s^2p)$ and $\text{seq}_2 = (p^2s^3)$, then $(\text{seq}_1 || \text{seq}_2) = (s^2pp^2s^3) = (s^2p^3s^3)$. For any $\text{seq} \in QS^*$, we define $(\text{seq} || \emptyset) = (\emptyset || \text{seq}) = \text{seq}$ and $(\text{seq} || \text{unbound}) = (\text{unbound} || \text{seq}) = \text{unbound}$.

3.1 Definition of Mixed CCA Security

Now we define mixed CCA security for a PKE scheme $\Pi = (\text{PKG}, \text{PEnc}, \text{PDec})$ as IND-ATK-like security parameterized by three query sequences $B, F, A \in QS^*$, denoted by $\langle B : F : A \rangle\text{-mCCA}$ security, via the $\langle B : F : A \rangle\text{-mCCA}$ experiment $\text{Expt}_{\Pi, \mathcal{A}}^{(B:F:A)\text{-mCCA}}(k)$ that an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ runs in:

$\text{Expt}_{\Pi, \mathcal{A}}^{(\mathbf{B}:\mathbf{F}:\mathbf{A})\text{-mCCA}}(k) : [(pk, sk) \leftarrow \text{PKG}(1^k); (m_0, m_1, \text{st}) \leftarrow \mathcal{A}_1^{\mathcal{O}}(pk); b \leftarrow \{0, 1\};$
 $c^* \leftarrow \text{PEnc}(pk, m_b); b' \leftarrow \mathcal{A}_2^{\mathcal{O}}(c^*, \text{st}) : \text{If } b' = b \text{ then return 1 else return 0}]$

where $\mathcal{O}(\cdot) = \text{PDec}(sk, \cdot)$ is a decryption oracle. However, how the decryption oracle is available is determined depending on the query sequences $\mathbf{B}, \mathbf{F}, \mathbf{A}$ in the following way: $\mathcal{A}_1$ can issue decryption queries following the sequence $\mathbf{B}$, and after all $\mathbf{B}$ -queries are completed, $\mathcal{A}_1$ can further issue decryption queries following the sequence $\mathbf{F}$. However, $\mathcal{A}_1$ need not complete all $\mathbf{F}$ -queries, and the ability to issue $\mathbf{F}$ -queries can be “shared” with $\mathcal{A}_2$. That is, as long as the order, the types, and the number of queries are maintained, the $\mathbf{F}$ -queries that $\mathcal{A}_1$ has not been completed can be taken over by $\mathcal{A}_2$. $\mathcal{A}_2$ can issue the remaining $\mathbf{F}$ -queries that $\mathcal{A}_1$ has left for $\mathcal{A}_2$, and after all $\mathbf{F}$ -queries are completed, $\mathcal{A}_2$ can further issue decryption queries following the sequence $\mathbf{A}$.² Moreover, as usual, $\mathcal{A}_2$ ’s queries must not contain the challenge ciphertext c^* .

We refer to $\mathbf{B}$, $\mathbf{F}$, and $\mathbf{A}$ as “**B**efore-challenge” queries, “**F**lexible” queries (in the sense that $\mathcal{A}$ can “flexibly” decide how it issues queries before/after the challenge), and “**A**fter-challenge” queries, respectively. For notational convenience, if $\mathbf{F} = \emptyset$ then we write $\langle \mathbf{B} :: \mathbf{A} \rangle\text{-mCCA}$, instead of $\langle \mathbf{B} : \emptyset : \mathbf{A} \rangle\text{-mCCA}$.

Definition 4. Let $\mathbf{B}, \mathbf{F}, \mathbf{A} \in \mathcal{QS}^*$. We say that a PKE scheme Π is $\langle \mathbf{B} : \mathbf{F} : \mathbf{A} \rangle\text{-mCCA}$ secure if $\text{Adv}_{\Pi, \mathcal{A}}^{(\mathbf{B}:\mathbf{F}:\mathbf{A})\text{-mCCA}}(k) = |\Pr[\text{Expt}_{\Pi, \mathcal{A}}^{(\mathbf{B}:\mathbf{F}:\mathbf{A})\text{-mCCA}}(k) = 1] - \frac{1}{2}|$ is negligible for any PPTA $\mathcal{A}$.

We define mixed CCA security for KEMs in exactly the same way as above.

With the mixed CCA security notions, we can express all the security notions mentioned in Section 2. These are summarized in Table 1. As noted earlier, for non-malleability, we adopt the characterization using a parallel query by Bellare and Sahai [3]. (In the table, we also include the *bounded parallel CCA* security notions defined below).

We remark that we can also define a parallel decryption query in mixed CCA security experiment (i.e. the $\langle \mathbf{B} : \mathbf{F} : \mathbf{A} \rangle\text{-mCCA}$ experiment) so that the number of ciphertexts contained in each parallel query is also bounded to be some predetermined value (say, t). However, such security definition is implied by $(|\langle \mathbf{B} || \mathbf{F} || \mathbf{A} \rangle| \cdot t)\text{-Bounded CCA}$ security, which is already achieved by the existing PKE schemes that are constructed only from IND-CPA secure schemes by the previous results [9,8]. Therefore, we think that studying security with such limitation is less interesting than studying mixed CCA security defined in this section, and is not treated in this paper.

Previously to this paper, Phan and Pointcheval [26] defined a similar notion which they call $(i, j)\text{-IND}$ security and $(i, j)\text{-NM}$ security, which are equivalent to $\langle \mathbf{s}^i :: \mathbf{s}^j \rangle\text{-mCCA}$ security and $\langle \mathbf{s}^i :: \mathbf{s}^j \mathbf{p} \rangle\text{-mCCA}$ security in our definition, respectively (for NM, we interpret it with parallel CCA-based characterization in [3]). They did not consider the “flexible” $\mathbf{F}$ -queries.

² In other words, in the $\langle \mathbf{B} : \mathbf{F} : \mathbf{A} \rangle\text{-mCCA}$ experiment, $\mathcal{A}_1$ can issue $(\mathbf{B} || \mathbf{F}_1)$ -queries, and $\mathcal{A}_2$ can issue $(\mathbf{F}_2 || \mathbf{A})$ -queries, for any pair of query sequences $(\mathbf{F}_1, \mathbf{F}_2)$ satisfying $(\mathbf{F}_1 || \mathbf{F}_2) = \mathbf{F}$, and how $\mathbf{F}$ is split into $\mathbf{F}_1$ and $\mathbf{F}_2$ can be decided adaptively by $\mathcal{A}$ in the experiment.

Table 1. Compatibility with Existing Security Notions

Existing Notions	Notation in Mixed CCA Security
IND-CPA	$\langle \emptyset :: \emptyset \rangle$ -mCCA
NM-CPA	$\langle \emptyset :: p \rangle$ -mCCA
IND- q -CCA2	$\langle \emptyset : s^q : \emptyset \rangle$ -mCCA
NM- q -CCA2	$\langle \emptyset : s^q : p \rangle$ -mCCA
IND- q -pCCA1	$\langle p^q :: \emptyset \rangle$ -mCCA
NM- q -pCCA1	$\langle p^q :: p \rangle$ -mCCA
IND- q -pCCA2	$\langle \emptyset : p^q : \emptyset \rangle$ -mCCA
NM- q -pCCA2	$\langle \emptyset : p^q : p \rangle$ -mCCA
IND-CCA1	$\langle \text{unbound} :: \emptyset \rangle$ -mCCA
NM-CCA1	$\langle \text{unbound} :: p \rangle$ -mCCA
IND-CCA2	$\langle \text{unbound} :: \text{unbound} \rangle$ -mCCA

Bounded Parallel CCA Security. Here, we define a natural and interesting special class of mixed CCA security which we call *bounded parallel CCA* security. This captures security against adversaries whose decryption queries are always parallel, and is a natural extension from the original bounded CCA security [9].

Depending on how the decryption oracle is available for an adversary, we define pCCA1 and pCCA2 as natural analogue of CCA1 and CCA2, respectively. Moreover, as is similar to the existing security notions, we define indistinguishability (IND) and non-malleability (NM).

Definition 5. Let $q \geq 0$ be an integer. We say that a PKE scheme is IND- q -pCCA1 (resp. IND- q -pCCA2, NM- q -pCCA1, and NM- q -pCCA2) secure if it is $\langle p^q :: \emptyset \rangle$ -mCCA (resp. $\langle \emptyset : p^q : \emptyset \rangle$ -mCCA, $\langle p^q :: p \rangle$ -mCCA, and $\langle \emptyset : p^q : p \rangle$ -mCCA) secure.

We define the bounded parallel CCA security notions for KEMs in the same way.

3.2 General Properties of Mixed CCA Security

Here, we show two general implication results about the mixed CCA security notions. (In this section, we always assume $B, F, A \in \mathcal{QS}^*$, and do not write it explicitly).

Firstly, by noticing the property of the “flexible” queries F , we obtain the following.

Theorem 1. For both PKE schemes and KEMs, $\langle B : F : A \rangle$ -mCCA security and “the combination of all security notions of the form $\langle (B||F_1) :: (F_2||A) \rangle$ -mCCA satisfying $(F_1||F_2) = F$ ” are equivalent.

The implication from the former to the latter is immediate by definition. Since the proof for the other direction is almost trivial, we omit the proof and only mention the intuition using the simplest case $F = s$. It is easy to see that $\langle B : s : A \rangle$ -mCCA adversary can be divided into two types: The first type that makes $\langle B||s \rangle$ -queries before the challenge, and A -queries after the challenge, and the

second type that makes B-queries before, and $(s||A)$ -queries after the challenge. Then, the experiment for the first type can be simulated by a $\langle(B||s) :: A\rangle$ -mCCA adversary while that for the second type can be simulated by a $\langle B :: (s||A)\rangle$ -mCCA adversary. This is easily extended to any $F \in \mathcal{QS}$ case. Note that if $F = \text{unbound}$, then the statement is again trivial because we can have $F_1 = F_2 = \text{unbound}$ (since $\text{unbound} = (\text{unbound}||\text{unbound})$), and thus in this case $\langle(B||F_1) :: (F_2||A)\rangle$ -mCCA security is equivalent to $\langle\text{unbound} :: \text{unbound}\rangle$ -mCCA = IND-CCA2 security, which implies all the mixed CCA security notions.

Next, we show that for PKE schemes with polynomially bounded plaintext space size and for KEMs, the A-queries, which is intended to be only available after the challenge, can actually be issued “flexibly” without destroying security.

Theorem 2. *For PKE schemes with polynomially bounded plaintext space size and for KEMs, $\langle B : F : A\rangle$ -mCCA security and $\langle B : (F||A) : \emptyset\rangle$ -mCCA security are equivalent.*

The implication from the latter notion to the former is immediate by definition. The proof for the other direction is given in the full version. Very roughly, showing the implication from the former notion to the latter is possible because the challenge ciphertext can be made “in advance” for PKE schemes with polynomially bounded plaintext space size and for KEMs. (In particular, for the PKE case, since the plaintext space size is polynomially bounded, the adversary’s two challenge plaintexts can be guessed correctly with probability $1/\text{poly}(k)$). Therefore, we can construct a reduction algorithm $\mathcal{B}$ that can successfully attack $\langle B : F : A\rangle$ -mCCA security using a successful $\langle B : (F||A) : \emptyset\rangle$ -mCCA adversary $\mathcal{A}$. Actually, in showing the proof, we have to be careful about the situation in which some of $\mathcal{A}$ ’s flexible decryption queries (i.e. $(F||A)$ -queries) issued by $\mathcal{A}$ before “ $\mathcal{A}$ ’s” challenge contains $\mathcal{B}$ ’s challenge ciphertext (which will be later used as $\mathcal{A}$ ’s challenge). However, the statistical property of PKE schemes and KEMs called *smoothness*, formalized in [2], guarantees that the probability of such a problematic situation occurring is negligible. For more details, we refer the reader to the full version.

4 Relations among Mixed CCA Security Notions

Due to its stage-specific queries and the difference between single and parallel queries, given two mixed CCA security notions, it is not always easy to tell if one notion implies the other. Therefore, a natural and yet non-trivial question is: *given two mixed CCA security notions $\langle B : F : A\rangle$ -mCCA and $\langle \tilde{B} : \tilde{F} : \tilde{A}\rangle$ -mCCA, under what conditions on $B, F, A, \tilde{B}, \tilde{F}, \tilde{A}$ are there implications/separations?*

In this section, we fully answer this question and show a necessary and sufficient condition for implications/separations between any two mixed CCA security notions. Interestingly, it turns out that *for PKE schemes, the relations among security notions are different depending on its plaintext space size*. The relations among mixed CCA security notions for PKE schemes with polynomially bounded plaintext space size and those for KEMs are always the same.

The rest of this section is organized as follows: In Section 4.1 we introduce a relation over query sequences which plays a key role for our results. Then in Sections 4.2 and 4.3 we show separation results and implication results, respectively. Finally in Section 4.4, we summarize the results by showing the necessary and sufficient conditions. For notational convenience, throughout this section we always assume $B, F, A, \tilde{B}, \tilde{F}, \tilde{A} \in \mathcal{QS}^*$. Due to space limitation, most of the proofs in this section are omitted and are given in the full version, and for theorems whose proofs are omitted, we provide some ideas for the proofs.

4.1 “is-Simulatable-by” Relation for Query Sequences

We first introduce the following relation over symbols.

Definition 6. We define a partial order “ $\subseteq_1$ ” over symbols $\{s, p\}$ by $s \subseteq_1 s$, $s \subseteq_1 p$, and $p \subseteq_1 p$.

Intuitively, the meaning of “ $\subseteq_1$ ” is that the former type oracle query “is-simulatable-by” the latter type of oracle query. The subscript “1” of “ $\subseteq_1$ ” denotes that it is a relation for one symbol, and it should not be mixed up with the relation for query sequences below (although the meaning is essentially the same).

Now, we extend the “is-simulatable-by” relation to query sequences:

Definition 7. Let $\text{seq}, \widetilde{\text{seq}} \in \mathcal{QS}^*$. We define a binary relation “ $\subseteq_{qs}$ ” over $\mathcal{QS}^*$ as follows. “ $\widetilde{\text{seq}} \subseteq_{qs} \text{seq}$ ” if and only if one of the following is satisfied:

- $\text{seq} = \text{unbound}$ or $\widetilde{\text{seq}} = \emptyset$
- $\text{seq} = (a_1 \dots a_m), \widetilde{\text{seq}} = (b_1 \dots b_n) \in \mathcal{QS} \setminus \{\emptyset\}$ where $a_i, b_j \in \{s, p\}$ for each $i \in [m], j \in [n]$, and there exists a strictly increasing function $f : [n] \rightarrow [m]$ such that $b_j \subseteq_1 a_{f(j)}$ holds for all $j \in [n]$.

If seq and $\widetilde{\text{seq}}$ do not satisfy the above, we write “ $\widetilde{\text{seq}} \not\subseteq_{qs} \text{seq}$ ”.

The subscript “qs” of $\subseteq_{qs}$ stands for *query sequence*. It is easy to see that the above relation “ $\subseteq_{qs}$ ” is a natural extension from $\subseteq_1$. Suppose $\widetilde{\text{seq}} \subseteq_{qs} \text{seq}$. Consider two PPTA adversaries $\mathcal{A}$ and $\mathcal{B}$ attacking a same PKE scheme, where $\mathcal{A}$ makes seq -queries and $\mathcal{B}$ makes $\widetilde{\text{seq}}$ -queries, and a situation in which $\mathcal{A}$ simulates the experiment for $\mathcal{B}$. If $\widetilde{\text{seq}} = \emptyset$, then $\mathcal{B}$ makes no query. If $\text{seq} = \text{unbound}$, then $\mathcal{A}$ can use unbounded oracle access, and thus $\mathcal{B}$ ’s decryption oracle can be simulated. Otherwise, (i.e. $\text{seq}, \widetilde{\text{seq}} \in \mathcal{QS} \setminus \{\emptyset\}$), then i -th query from $\mathcal{B}$ can be simulated by $\mathcal{A}$ ’s $f(i)$ -th query (where f is a strictly increasing function guaranteed to exist by definition) for all $i \in [n]$.

Now, given two sequences $\text{seq}, \widetilde{\text{seq}} \in \mathcal{QS}^*$ we can tell if $\widetilde{\text{seq}} \subseteq_{qs} \text{seq}$ or $\widetilde{\text{seq}} \not\subseteq_{qs} \text{seq}$.³ For example, $(s^2p) \subseteq_{qs} (psps)$; $(sp^2) \not\subseteq_{qs} (s^2ps^2)$; $s^r \subseteq_{qs} s^q$ iff $q \geq r$.

³ Note that “ $\subseteq_{qs}$ ” forms a partial order over $\mathcal{QS}^*$. However, it is not a total order. For example, we have both $(sp) \not\subseteq_{qs} (ps)$ and $(ps) \not\subseteq_{qs} (sp)$.

4.2 Separation Results

A common approach for showing a separation of a security notion $\mathbf{X}$ from a security notion $\mathbf{Y}$ for PKE schemes is to construct a “separating” PKE scheme from a building block $\mathbf{X}$ -secure PKE scheme: the decryption algorithm of the separating scheme typically has some “backdoor” mechanism, which leads to some “critical information” v (e.g. secret key for the building block PKE scheme) for breaking $\mathbf{Y}$ security so that $\mathbf{Y}$ -adversary can, by using a decryption oracle, reach for v and break $\mathbf{Y}$ -security of the separating PKE scheme while an $\mathbf{X}$ -adversary cannot reach for v or simply v is useless for breaking $\mathbf{X}$ -security of the scheme. We also follow this approach.

Useful Tool for Separation: Backdoor-Sequence Scheme. In order for the above approach to work, what to use as the critical information and how to implement such backdoor mechanism are the main issues. We wish to implement a backdoor mechanism so that given two sequences $\text{seq}, \widetilde{\text{seq}} \in \mathcal{QS}^*$, if $\widetilde{\text{seq}} \not\subseteq_{qs} \text{seq}$,⁴ then an adversary making $\widetilde{\text{seq}}$ -queries can finally reach for a critical information (and break some security of a separating PKE scheme) while an adversary making seq -queries cannot. This is indeed possible. We can implement such backdoor mechanism as a *sequence of backdoor information* $(u_1, \dots, u_{|\widetilde{\text{seq}}|+1})$ and a *strategy for “how to release next backdoor information”*, based on $\text{seq}, \widetilde{\text{seq}}$, and the critical information v . Specifically, let $\widetilde{\text{seq}} = (b_1 \dots b_n)$ such that $b_i \in \{\mathbf{s}, \mathbf{p}\}$ for $i \in [n]$.

- The sequence of backdoor information $(u_1, \dots, u_{n+1})$ is set up so that $u_1 = 1^k$ (any publicly known value will do), $u_2, \dots, u_n$ are random values (which must be hard to guess), and u_{n+1} is the critical information v .
- The strategy for “how to release next backdoor information”, depending on $\widetilde{\text{seq}} = (b_1 \dots b_n)$, is set up so that: If $b_i = \mathbf{s}$, this “release strategy” on input u_i outputs u_{i+1} directly; If $b_i = \mathbf{p}$, this “release strategy” on input u_i (together with some index j) outputs a (j -th) “secret-share” of u_{i+1} , so that if we collect the shares more than a threshold which is set to be a value greater than $|\text{seq}|$, we can reconstruct u_{i+1} .

(Intuitively, such a release strategy is implemented into a decryption algorithm of a separating PKE scheme so that if the decryption algorithm takes some special information indicating “backdoor mode” as input, the output of the release strategy is used instead of decrypting as a ciphertext). Constructed as above, an adversary making $\widetilde{\text{seq}}$ -queries to the release strategy can finally obtain $u_{n+1} = v$. In particular, if $b_i = \mathbf{p}$ then an adversary can make a parallel query to the release strategy to obtain all the share of u_{i+1} at once, and thus can reconstruct u_{i+1} . It is actually possible to show that *if $\widetilde{\text{seq}} \not\subseteq_{qs} \text{seq}$, then no adversary who is only allowed to make seq -queries to the release strategy can reach for $u_{n+1} = v$* , and thus we can make a difference in the information available for an adversary making $\widetilde{\text{seq}}$ -queries and that making seq -queries.

⁴ As we have seen in Section 4.1, if $\widetilde{\text{seq}} \subseteq_{qs} \text{seq}$, then any information available for adversaries making $\widetilde{\text{seq}}$ -queries is also available for those making seq -queries.

In order to make it easier to analyze PKE schemes used to show separations, in the full version we formalize this “backdoor mechanism” as a “stand alone” primitive. We call it a *backdoor-sequence scheme*, and use it as one of main building blocks for constructing the separating schemes that are used to establish the separations in the following paragraphs in this subsection.

Separation by Total Query Sequence.

Theorem 3. *For both PKE schemes and KEMs, if $(\tilde{B}||\tilde{F}||\tilde{A}) \not\prec_{qs} (B||F||A)$, then $\langle B : F : A \rangle$ -mCCA security does not imply $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA security.*

The idea for building the separating PKE scheme for showing Theorem 3 is straightforward. We use the secret key sk for the building block scheme of the separating PKE scheme as a critical information, and setup the backdoor-sequence scheme appropriately. That is, a $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA adversary who can (in total) make $(\tilde{B}||\tilde{F}||\tilde{A})$ -queries can finally reach for sk and decrypt the challenge ciphertext. However, since $(\tilde{B}||\tilde{F}||\tilde{A}) \not\prec_{qs} (B||F||A)$, the property of the backdoor-sequence scheme guarantees that a $\langle B : F : A \rangle$ -mCCA adversary who is only allowed to make $(B||F||A)$ -queries in total cannot reach for it, and thus the separating PKE scheme remains $\langle B : F : A \rangle$ -mCCA secure. The same proof strategy works for the KEM case.

Separation by After-challenge Queries.

Theorem 4. *For both PKE schemes and KEMs, if $(\tilde{F}||\tilde{A}) \not\prec_{qs} (F||A)$, then $\langle B : F : A \rangle$ -mCCA security does not imply $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA security.*

For an explanation here, consider the extreme case: the separation of $\langle \emptyset : \tilde{A} \rangle$ -mCCA security from $\langle \text{unbound} : A \rangle$ -mCCA security under the condition $\tilde{A} \not\prec_{qs} A$. Note that a $\langle \text{unbound} : A \rangle$ -mCCA adversary can make unbounded single queries before the challenge while a $\langle \emptyset : \tilde{A} \rangle$ -mCCA adversary can make no query. Therefore, the critical information for breaking $\langle \emptyset : \tilde{A} \rangle$ -mCCA security must be something that is useful and available only after the challenge. We set the critical information to be the decryption result of a ciphertext (which can be a challenge ciphertext) that is input together with the backdoor information into the decryption algorithm of the separating scheme, and use a pseudorandom function F to realize the separating PKE scheme that has a “ciphertext-dependent” backdoor sequence. More specifically, a seed K for F is picked as a part of a secret key of the separating PKE scheme. The decryption algorithm of the separating scheme, on input a ciphertext c together with backdoor and some information that indicates “backdoor mode”, derives a pseudorandom value $R = F_K(c)$ and use this R as a randomness for deriving the sequence of backdoors, and then outputs a corresponding “next backdoor”. Since the backdoor-sequence scheme is set up so that an adversary that can make $\tilde{A}$ -queries can finally reach for the critical information (decryption of any ciphertext), a $\langle \emptyset : \tilde{A} \rangle$ -mCCA adversary can finally reach for the decryption result of the challenge ciphertext by appropriately making decryption queries after the challenge. However, since $\tilde{A} \not\prec_{qs} A$, the same does not apply to a $\langle \text{unbound} : A \rangle$ -mCCA adversary that can make only

A-queries after the challenge. Therefore, to break $\langle \text{unbound} :: A \rangle$ -mCCA security of the separating scheme, the adversary has to essentially break $\langle \text{unbound} :: A \rangle$ -mCCA security of the building block scheme (unless it breaks security of the pseudorandom function or the backdoor-sequence scheme). Essentially the same proof strategy works for proving the KEM case. Using a pseudorandom function to set up “ciphertext-dependent” backdoor information was previously used in [1] to separate $\text{NM-CCA2} = \text{IND-CCA2}$ security from NM-CCA1 security.

Separation by Before-challenge Queries.

Theorem 5. *For PKE schemes with superpolynomially large plaintext space size, if $\langle \tilde{B} || F \rangle \not\subseteq_{qs} (B || F)$, then $\langle B : F : A \rangle$ -mCCA security does not imply $\langle \tilde{B} : \tilde{F} : A \rangle$ -mCCA security.*

Note that this theorem is only true for PKE schemes with superpolynomially large plaintext space size. For an explanation here, consider the extreme case: the separation of $\langle \tilde{B} :: \emptyset \rangle$ -mCCA security from $\langle B :: \text{unbound} \rangle$ -mCCA security under the condition $\tilde{B} \not\subseteq_{qs} B$. This time, the critical information for breaking $\langle \tilde{B} :: \emptyset \rangle$ -mCCA security must be something that is useful only before the challenge, because $\langle \tilde{B} :: \emptyset \rangle$ -mCCA adversary can make no query after the challenge. We use a one-way function f to construct the separating PKE scheme so that it has “weak” plaintexts, which are not encrypted at all by the encryption algorithm of the separating PKE scheme. (Similar ideas are used in [26,24,2]). A public key of the separating PKE scheme contains $V = f(m^*)$ for some random element m^* chosen from the plaintext space of the underlying PKE scheme, and weak plaintexts m are the ones satisfying $f(m) = V$. We set the critical information to be m^* itself, where the backdoor-sequence scheme will finally release m^* if $\tilde{B}$ -queries are appropriately performed, and thus m^* can be used as one of two challenge plaintexts to break $\langle \tilde{B} :: \emptyset \rangle$ -mCCA security. However, since $\tilde{B} \not\subseteq_{qs} B$, the property of the backdoor-sequence scheme guarantees that a $\langle B :: \text{unbound} \rangle$ -mCCA adversary cannot reach for m^* before the challenge (unless it break one-wayness of f). Moreover, the weak plaintext m^* is useless even if it is found after the challenge. Therefore, in order to break $\langle B :: \text{unbound} \rangle$ -mCCA security of the separating scheme, the adversary essentially has to attack $\langle B :: \text{unbound} \rangle$ -mCCA security of the building block scheme.

4.3 Implication Results

A combination of Theorems 3, 4, and 5 shows that given two mixed CCA security notions $\langle B : F : A \rangle$ -mCCA and $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA, the latter notion is separated from the former if $\langle \tilde{B} || \tilde{F} || \tilde{A} \rangle \not\subseteq_{qs} (B || F || A)$, $\langle \tilde{B} || \tilde{F} \rangle \not\subseteq_{qs} (B || F)$, or $\langle \tilde{F} || \tilde{A} \rangle \not\subseteq_{qs} (F || A)$ holds for PKE schemes with superpolynomially large plaintext space. We show that if none of the above conditions are satisfied, then we actually have an implication from the former notion to the latter, where this implication is also true for all PKE schemes and KEMs.

Theorem 6. *For both PKE schemes and KEMs, if $\langle \tilde{B} || \tilde{F} || \tilde{A} \rangle \subseteq_{qs} (B || F || A)$, $\langle \tilde{B} || \tilde{F} \rangle \subseteq_{qs} (B || F)$, and $\langle \tilde{F} || \tilde{A} \rangle \subseteq_{qs} (F || A)$ hold simultaneously, then $\langle B : F : A \rangle$ -mCCA security implies $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA security.*

This theorem holds because it can be shown that if the three conditions regarding query sequences are satisfied, then whatever strategy regarding the “flexible” queries an $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA adversary may take, the $\langle B : F : A \rangle$ -mCCA experiment can be perfectly simulated by an $\langle B : F : A \rangle$ -mCCA adversary.

Combining Theorem 6 with Theorem 2, we obtain the following corollary.

Corollary 1. *For PKE schemes with polynomially bounded plaintext space size and for KEMs, if $(\tilde{B}||\tilde{F}||\tilde{A}) \subseteq_{qs} (B||F||A)$ and $(\tilde{F}||\tilde{A}) \subseteq_{qs} (F||A)$ hold simultaneously, then $\langle B : F : A \rangle$ -mCCA security implies $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA security.*

Proof. By Theorem 2, we know that for PKE schemes with polynomially bounded plaintext space size and for KEMs, $\langle B : F : A \rangle$ -mCCA security implies $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA security if and only if $\langle B : (F||A) : \emptyset \rangle$ -mCCA security implies $\langle \tilde{B} : (\tilde{F}||\tilde{A}) : \emptyset \rangle$ -mCCA security. Then, Theorem 6 tells us that the sufficient condition of the implication from the former notion to the latter is: “ $(\tilde{B}||(\tilde{F}||\tilde{A})||\emptyset) \subseteq_{qs} (B||(\tilde{F}||\tilde{A})||\emptyset)$, $(\tilde{B}||(\tilde{F}||\tilde{A})) \subseteq_{qs} (B||(\tilde{F}||\tilde{A}))$, and $((\tilde{F}||\tilde{A})||\emptyset) \subseteq_{qs} ((F||A)||\emptyset)$ hold simultaneously.” Simplifying this condition yields Corollary 1. $\square$

4.4 Necessary and Sufficient Conditions for Implication/Separation

As a summarization of the results in this section, we show the following necessary and sufficient conditions for implication/separation among mixed CCA security.

Theorem 7. *For PKE schemes with superpolynomially large plaintext space size, $\langle B : F : A \rangle$ -mCCA security implies $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA security if and only if $(\tilde{B}||\tilde{F}||\tilde{A}) \subseteq_{qs} (B||F||A)$, $(\tilde{B}||\tilde{F}) \subseteq_{qs} (B||F)$, and $(\tilde{F}||\tilde{A}) \subseteq_{qs} (F||A)$ hold simultaneously.*

Proof. This follows from a combination of Theorems 3, 4, 5, and 6. $\square$

Theorem 8. *For PKE schemes with polynomially bounded plaintext space size and for KEMs, $\langle B : F : A \rangle$ -mCCA security implies $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA security if and only if $(\tilde{B}||\tilde{F}||\tilde{A}) \subseteq_{qs} (B||F||A)$ and $(\tilde{F}||\tilde{A}) \subseteq_{qs} (F||A)$ hold simultaneously.*

Proof. This follows from a combination of Theorems 3, 4, and Corollary 1. $\square$

We believe the relations among security notions shown in this section are useful for future studies on PKE schemes and KEMs whose security notions can be expressed in mixed CCA security notions. For example, by utilizing the above theorems, we can fully establish the relations among bounded parallel CCA security and other existing security notions in Table 1. We also note that the previously established relations among security notions [1,9,16] can be re-proved as corollaries from the above theorems.

Importance of Plaintext Space Size in Relations among Security Notions for PKE Schemes. As our results in this section have clarified, it is important to care about the size of the plaintext space size for relations among security notions for PKE schemes. Specifically, Theorems 7 and 8 tell us that given $\langle B : F : A \rangle$ -mCCA and $\langle \tilde{B} : \tilde{F} : \tilde{A} \rangle$ -mCCA security notions, the implication/separation from the

former notion to the latter notion differs if $(\tilde{B}||\tilde{F}||\tilde{A}) \subseteq_{qs} (B||F||A)$, $(\tilde{F}||\tilde{A}) \subseteq_{qs} (F||A)$, and $(\tilde{B}||\tilde{F}) \not\subseteq_{qs} (B||F)$ hold simultaneously.

5 Feasibility Results from IND-CPA Secure PKE Schemes

By adopting the notion of mixed CCA security, in this section we show two black-box constructions of PKE schemes, which can encrypt plaintexts of polynomial length (thus, exponentially large plaintext space), from IND-CPA secure schemes.

The first result is the following.

Theorem 9. *For any polynomial $q \geq 0$, there exists a shielding black-box construction of a $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA secure PKE scheme which can encrypt plaintexts of polynomial length from an IND-CPA secure PKE scheme.*

Proof. This theorem is proved by combining the existing results and Theorem 2 in Section 3.2. The following statement is due to the result by Choi et al. [8].

Lemma 1. [8] *For any polynomial $q \geq 0$, there exists a shielding black-box construction of an NM- q -CCA2 secure PKE scheme which can encrypt plaintexts of polynomial length from an IND-CPA secure PKE scheme.*

Recall that $\text{NM-}q\text{-CCA2} = \langle \emptyset : s^q : p \rangle$ -mCCA (see Table 1). Since any $\langle B : F : A \rangle$ -mCCA secure PKE scheme can be trivially used as a KEM with the same security by encrypting a uniformly random string K and using it as a session-key, Lemma 1 implies that we can construct a $\langle \emptyset : s^q : p \rangle$ -mCCA secure KEM (we call it the *CDMW KEM*) from any IND-CPA secure PKE scheme in a black-box and shielding manner. Then, by Theorem 2 for KEMs, we can immediately say that the CDMW KEM is $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA secure. Finally, by combining the CDMW KEM with an IND-CCA2 secure DEM, we obtain the desired result. (It is implicit from the works by Cramer and Shoup [10] and by Herranz et al. [16] that if we combine a $\langle B : F : A \rangle$ -mCCA secure KEM and an IND-CCA2 secure DEM in a straightforward manner, we can obtain a $\langle B : F : A \rangle$ -mCCA secure PKE scheme). Note that we can construct an IND-CCA2 secure DEM even without any computational assumption (see e.g. [10, Section 7.2.2]). Moreover, the “shielding” and “black-box” properties are trivially preserved by our construction. This completes the proof of Theorem 9. $\square$

$\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA security implies $\text{NM-}q\text{-CCA2} = \langle \emptyset : s^q : p \rangle$ -mCCA security by definition, while by Theorem 7 we know that for PKE schemes with superpolynomially large plaintext space size, $\langle \emptyset : s^q : p \rangle$ -mCCA security does not imply $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA security. Therefore, for these types of PKE schemes, $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA security is strictly stronger than $\text{NM-}q\text{-CCA2}$ security.

We remark that Theorem 2 actually implies that the original CDMW PKE scheme [8] already achieves the shielding black-box construction of $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA secure PKE schemes *if it is used with short plaintexts (so that the plaintext space size is bounded to be polynomial)*. However, the Choi et al. result itself does not imply Theorem 9, because it is not obvious how to construct $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA secure PKE schemes which can encrypt plaintexts of polynomially length

from PKE schemes that satisfies the same security but has only polynomially bounded plaintext space size, in a black-box and shielding manner⁵.

We also remark that the original CDMW PKE scheme [8] might be shown to be $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA secure as it is for large plaintext space size, using the same assumptions used to show its NM- q -CCA2 security. However, our main purpose here is to show the improved feasibility rather than the concrete construction and efficiency, and thus we did not try proving directly that the CDMW PKE is $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA secure.

Theorem 9 implies the following corollary.

Corollary 2. *There exists a shielding black-box construction of an IND-1-pCCA2 secure PKE scheme which can encrypt plaintexts of polynomial length from an IND-CPA secure PKE scheme.*

Our second result on black-box constructions is the following.

Theorem 10. *For any polynomials $q, q' \geq 0$, there exists a shielding black-box construction of a $\langle s^q p : s^{q'} : \emptyset \rangle$ -mCCA secure PKE scheme which can encrypt plaintexts of polynomial length from an IND-CPA secure PKE scheme.*

Proof. To prove this theorem, we will use the following result which is implicit from [9, Lemma 1]⁶:

Lemma 2. *(Implicit from [9]). For any $B \in \mathcal{QS}^*$ and any polynomial $q' \geq 0$, there exists a shielding black-box construction of a $\langle B : s^{q'} : \emptyset \rangle$ -mCCA secure PKE scheme from a $\langle B :: \emptyset \rangle$ -mCCA secure PKE scheme.*

We call the construction by Cramer et al. [9] the *CHH+ PKE scheme*. Due to Theorem 9 above, for any polynomial $q \geq 0$, we can construct a $\langle \emptyset : s^q p : \emptyset \rangle$ -mCCA secure PKE scheme, which is also $\langle s^q p :: \emptyset \rangle$ -mCCA secure, from any IND-CPA secure PKE scheme. Then, by using this PKE scheme as a building block of the CHH+ PKE scheme, due to Lemma 2, we have a PKE scheme which satisfies the claimed security. The CHH+ PKE construction is shielding and black-box. Since the construction of the PKE scheme in Theorem 9 is also shielding and black-box, so is the construction as a whole. The size of the plaintext space is maintained as well. This completes the proof of Theorem 10. $\square$

We note that by Theorem 7, for PKE schemes with superpolynomially large plaintext space size, $\langle s^q p : s^{q'} : \emptyset \rangle$ -mCCA security achieved in Theorem 10 cannot be directly compared with the notion achieved in Theorem 9 (actually even with NM-CPA = $\langle \emptyset :: p \rangle$ -mCCA security). However, the security achieved in Theorem 10

⁵ Recently, Myers and Shelat [21] showed a black-box construction of multi-bit IND-CCA2 secure PKE schemes from 1-bit IND-CCA2 secure PKE schemes. Though it seems that their result can be extended (with some modification) to any mixed CCA security, we remark that their construction is non-shielding.

⁶ The original statement of Lemma 1 in [9] shows a special case of Lemma 2 in which $B = \emptyset$. Moreover, the special case of Lemma 2 in which $B = \text{unbound}$ is also mentioned in [9]. See Remark 2 after the proof of Lemma 1 in [9].

allows the bounded number of “flexible” single queries before and after the challenge, after the parallel query in the first stage, while the security achieved by Theorem 9 does not allow any query after one parallel query for an adversary. Thus we believe that Theorem 10 is also interesting as a feasibility result.

Handling Decryption of Unboundedly Many Ciphertexts before the Challenge. Previous to our work, none of the constructions of PKE schemes that use only IND-CPA secure ones have achieved security against adversaries that can observe unboundedly many decryption results (via the decryption oracle) in the first stage, i.e., before choosing two challenge plaintexts, regardless of whether the construction is black-box or non-black-box. On the other hand, the constructions in Theorems 9 and 10 (and also the combination of [8] and Theorem 2) achieve security against adversaries that can observe unboundedly many decryption results by one parallel decryption query before the challenge.

Thus, due to the results in this section, it has been clarified that *the difficulty of constructing an IND-CCA1 secure PKE scheme only from IND-CPA secure ones lies not in whether the number of decryption results that the adversary can see before the challenge is bounded or not, but in whether the number of the adversary’s “adaptive” decryption queries is bounded.*

6 Open Problems

Constructions Secure against Two or More Parallel Queries. None of our feasibility results achieves mixed CCA security in which we can handle more than one parallel decryption query, and whether we can construct a PKE scheme with such security only using IND-CPA secure schemes is still unclear. Therefore, we would like to leave it as an open problem. Since any (unbounded) CCA secure PKE construction from IND-CPA secure ones must first be secure against adversaries who make two or more parallel decryption queries, we believe that overcoming this barrier of “two parallel queries” is worth tackling.

We notice that if we can, by only using an IND-CPA secure PKE scheme as a building block, construct a (strong) *designated verifier* (DV) NIZK proof system [23,9] for any NP language with q -bounded “parallel” strong soundness, which is a natural extension of a (strong) DV-NIZK with q -bounded strong soundness [9] in the soundness experiment of which an adversary can ask verification of many theorem/proof pairs in a parallel manner, then by using the DV-NIZK proof system in the Dolev-Dwork-Naor construction [11,23,9] (resp. the Naor-Yung construction [22]) we will be able to construct an IND- $(q + 1)$ -pCCA2 (resp. IND- $(q + 1)$ -pCCA1) secure PKE scheme. However, how to construct such a DV-NIZK proof system only from IND-CPA secure PKE schemes is not known so far. This might be worth looking at towards the next step from our results.

Stronger Black-Box Impossibility Results. Since the constructions in Theorems 9 and 10 are shielding and black-box, according to the impossibility result of [12] and the transitivity of black-box constructions, we have that there exists no shielding black-box construction of an IND-CCA1 secure PKE scheme from PKE schemes which satisfy any security notion achieved in Theorems 9 and 10.

It would also be interesting to clarify if we can show a stronger impossibility result than [12] such that constructing IND- q -pCCA1 secure PKE schemes from IND-CPA secure one in a shielding and black-box manner for some $q > 1$ is impossible. (Or more generally, we can also consider the impossibility of some of mixed CCA security notion). Note that this strengthening of the impossibility result of [12] can make sense only if we consider parallel decryption queries, because the result by Choi et al. [8] already shows that it is possible to achieve the strongest form of (ordinary) bounded CCA security, namely, NM- q -CCA2, in a black-box and shielding manner.

Acknowledgement

The authors would like to thank Jacob Schuldt for his helpful comments and suggestions. The authors also would like to thank anonymous reviewers of PKC 2011 for their invaluable comments.

References

1. 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)
2. Bellare, M., Hofheinz, D., Kiltz, E.: Subtleties in the definition of IND-CCA: When and how should challenge-decryption be disallowed?, Cryptology ePrint Archive: Report 2009/418 (2009)
3. 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)
4. Bellare, M., Yung, M.: Certifying permutations: Noninteractive zero-knowledge based on any trapdoor permutation. J. Cryptology 9(3), 149–166 (1996)
5. Boneh, D., Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. SIAM J. Comput. 36(5), 1301–1328 (2007)
6. Boyen, X., Mei, Q., Waters, B.: Direct chosen ciphertext security from identity-based techniques. In: CCS 2005, pp. 320–329 (2005)
7. 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)
8. Choi, S.G., Dachman-Soled, D., Malkin, T., Wee, H.: Black-box construction of a non-malleable encryption scheme from any semantically secure one. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 427–444. Springer, Heidelberg (2008)
9. 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)
10. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM J. Comput. 33(1), 167–226 (2003)
11. Dolev, D., Dwork, C., Naor, M.: Non-malleable cryptography. SIAM J. Comput. 30(2), 391–437 (2000)

12. Gertner, Y., Malkin, T., Myers, S.: Towards a separation of semantic and CCA security for public key encryption. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 434–455. Springer, Heidelberg (2007)
13. Goldwasser, S., Micali, S.: Probabilistic encryption. *J. Comput. Syst. Sci.* 28(2), 270–299 (1984)
14. 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)
15. 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)
16. Herranz, J., Hofheinz, D., Kiltz, E.: Some (in)sufficient conditions for secure hybrid encryption. *Inf. Comput.* 208(11), 1243–1257 (2010)
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. Hofheinz, D., Kiltz, E.: Practical chosen ciphertext secure encryption from factoring. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 313–332. Springer, Heidelberg (2009)
19. Kiltz, E., Mohassel, P., O’Neill, A.: Adaptive trapdoor functions and chosen-ciphertext security. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 673–692. Springer, Heidelberg (2010)
20. 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)
21. Myers, S., Shelat, A.: Bit encryption is complete. In: FOCS 2009, pp. 607–616 (2009)
22. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: STOC 1990, pp. 427–437 (1990)
23. Pass, R., Shelat, A., Vaikuntanathan, V.: Construction of a Non-malleable Encryption Scheme from Any Semantically Secure One. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 271–289. Springer, Heidelberg (2006)
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. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: STOC 2008, pp. 187–196 (2008)
26. Phan, D.H., Pointcheval, D.: On the security notions for public-key encryption schemes. In: Blundo, C., Cimato, S. (eds.) SCN 2004. LNCS, vol. 3352, pp. 33–46. Springer, Heidelberg (2005)
27. 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)
28. Reingold, O., Trevisan, L., Vadhan, S.: Notions of reducibility between cryptographic primitives. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 1–20. Springer, Heidelberg (2004)
29. Rosen, A., Segev, G.: Chosen-ciphertext security via correlated products. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 419–436. Springer, Heidelberg (2009)
30. Wee, H.: Efficient chosen-ciphertext security via extractable hash proofs. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 314–332. Springer, Heidelberg (2010)