

Trading Plaintext-Awareness for Simulatability to Achieve Chosen Ciphertext Security

Takahiro Matsuda and Goichiro Hanaoka

National Institute of Advanced Industrial Science and Technology (AIST), Tokyo,
Japan {t-matsuda,hanaoka-goichiro}@aist.go.jp

Abstract. In PKC 2014, Dachman-Soled showed a construction of a chosen ciphertext (CCA) secure public key encryption (PKE) scheme based on a PKE scheme which simultaneously satisfies a security property called weak simulatability and (standard model) plaintext awareness (sPA1) in the presence of multiple public keys. It is not well-known if plaintext awareness for the multiple keys setting is equivalent to the more familiar notion of that in the single key setting, and it is typically considered that plaintext awareness is a strong security assumption (because to achieve it we have to rely on a “knowledge”-type assumption). In Dachman-Soled’s construction, the underlying PKE scheme needs to be plaintext aware in the presence of $2k + 2$ public keys.

The main result in this work is to show that the strength of plaintext awareness required in the Dachman-Soled construction can be somehow “traded” with the strength of a “simulatability” property of other building blocks. Furthermore, we also show that we can “separate” the assumption that a single PKE scheme needs to be both weakly simulatable and plaintext aware in her construction. Specifically, in this paper we show two new constructions of CCA secure key encapsulation mechanisms (KEMs): Our first scheme is based on a KEM which is chosen plaintext (CPA) secure and plaintext aware only under the 2 keys setting, and a PKE scheme satisfying a “slightly stronger” simulatability than weak simulatability, called “trapdoor simulatability” (introduced by Choi et al. ASIACRYPT 2009). Our second scheme is based on a KEM which is 1-bounded CCA secure (Cramer et al. ASIACRYPT 2007) and plaintext aware only in the *single* key setting, and a trapdoor simulatable PKE scheme. Our results add new recipes for constructing CCA secure PKE/KEM from general assumptions (that are incomparable to those used by Dachman-Soled), and in particular show interesting trade-offs among building blocks with those used in Dachman-Soled’s construction.

Keywords: public key encryption, key encapsulation mechanism, chosen ciphertext security, plaintext-awareness, trapdoor simulatability.

1 Introduction

1.1 Background and Motivation

For public key encryption (PKE), security (indistinguishability) against chosen ciphertext attacks (CCA) [46, 49, 23] is nowadays considered as a de-facto

standard security notion required in most practical situations/applications in which PKE schemes are used. CCA security is quite important in both practical and theoretical points of view. It implies security against practical attacks (e.g. Bleichenbacher’s attack [8]) and it also implies very strong and useful security notions, such as non-malleability [23] and universal composability [10]. Thus, constructing and understanding CCA secure PKE schemes is one of the central research themes in the area of cryptography. In this paper, we focus on the constructions of CCA secure PKE schemes and its closely related primitive called *key encapsulation mechanism* (KEM) from general cryptographic assumptions. There have been a number of works that show that several different kinds of cryptographic primitives are sufficient to realize CCA secure PKE/KEM: These include trapdoor permutations [23] (with some enhanced property [27]), identity-based encryption [12] and a weaker primitive called tag-based encryption [32], lossy trapdoor function [48] and related primitives [50, 42, 33, 54, 13], PKE schemes with weaker-than-but-close-to-CCA security [31, 34, 40], positive results on cryptographic obfuscation [52, 38], the combination of a CPA secure PKE scheme and a strong form of hash functions [39], and very recently, the combination of a sender non-committing encryption scheme and a key-dependent-message secure symmetric key encryption (SKE) scheme [41]. (We review more works in Section 1.4.)

In PKC 2014, Dachman-Soled [18] showed a construction of a CCA secure PKE scheme based on a PKE scheme which simultaneously satisfies a security property called weak simulatability [20, 43] and (standard model) plaintext awareness (sPA1) [5] in the presence of multiple public keys [43], which is based on the earlier work by Myers, Sergi, and Shelat [43] who showed a construction of a PKE scheme that achieves security slightly weaker than CCA (the so-called cNM-CCA1 security). Plaintext awareness was first introduced by Bellare and Rogaway [7] as a useful notion for showing CCA security of a PKE scheme in the random oracle model [6], and was used in a number of random-oracle-model constructions (e.g. [7, 24, 25, 47]). Bellare and Palacio [5] defined the standard model versions of plaintext awareness.¹ The plaintext awareness notions were further studied by subsequent works (e.g. [20]). The most works on plaintext awareness studied the notions for the single key setting. The extension to the multiple keys setting was first introduced by Myers, Sergi, and Shelat [43].

We note that it is not well-known or well-studied if plaintext awareness for the multiple keys setting is equivalent to the more familiar notion of plaintext awareness in the single key setting, and it is typically considered that plaintext awareness is a strong security assumption (because to achieve it we have to rely on a “knowledge”-type assumption). In the construction of [18], the underlying PKE scheme needs to be plaintext aware in the presence of $2k + 2$ public keys. Our motivation in this work is to clarify whether we can weaken the assumption

¹ [5] defined several versions (PA0, PA1, and PA2, with their computational/statistical/perfect variants) for standard model plaintext awareness. As in the previous works [43, 18], we focus on the statistical PA1 notion in the multiple keys setting (denoted by “sPA1 $_{\ell}$ ”, where ℓ denotes the number of public keys).

of plaintext awareness in Dachman-Soled’s construction [18]. As mentioned in [18], a plaintext aware (sPA1) PKE scheme seems almost like a CCA1 secure PKE scheme [46], but it seems not possible to replace the building block PKE scheme in [18] with a CCA1 secure scheme to remove the plaintext awareness. It is currently not known if we can construct a CCA secure PKE scheme only from a CPA secure scheme or even from a CCA1 secure scheme. We believe that studying the possibility of weakening the assumption of plaintext awareness from [18] thus is expected to lead to deepening our knowledge on this topic, and generally contribute to the long line of research on clarifying the minimal general assumption that implies CCA secure PKE.

1.2 Our Contributions

Based on the motivation mentioned above, we study the possibility of weakening the requirements of plaintext awareness used in Dachman-Soled’s construction [18], and come up with new results that show that the strength of plaintext awareness required in [18] can be somehow “traded” with the strength of a “simulatability” property of other building blocks. Furthermore, we also show that we can “separate” the requirement that a single PKE scheme needs to be simultaneously weakly simulatable and plaintext aware, in her construction.

Specifically, in this paper we show two new constructions of CCA secure KEMs (which are given in Section 4), based on the assumptions that are *incomparable* to those used in [18]:

- Our first construction (Section 4.1) is based on a KEM which is chosen plaintext (CPA) secure and plaintext aware only under the 2 keys setting², and a PKE scheme satisfying a “slightly stronger” simulatability than weak simulatability, called “trapdoor simulatability” (introduced by Choi et al. [14]). Actually, although we write that it is “slightly stronger”, it is formally *incomparable* to weak simulatability. For more details, see Section 1.3.
- Our second construction (Section 4.2) is based on a KEM which is 1-bounded CCA secure [15] and plaintext aware only in the *single* key setting, and a trapdoor simulatable PKE scheme. We can in fact slightly weaken the requirement of 1-bounded CCA security to CPA security in the presence of one “plaintext-checking” query [47, 1]. We will also show that we can construct a KEM satisfying simultaneously 1-bounded CCA security and plaintext awareness under the single key setting, based on a KEM satisfying CPA security and plaintext awareness under the $2k$ keys setting, via the recent result by Dodis and Fiore [21, Appendix C].

One may wonder the meaning of the second construction, because if we use a KEM that is plaintext aware under $O(k)$ keys setting, there is no merit compared to our first construction. We are however considering it to be still meaningful in several aspects, and we refer the reader to Section 4.2 for more discussions regarding the second construction.

² Plaintext awareness for KEMs is defined analogously to that for PKE. See Section 2.1.

Note that from CCA secure KEMs, we can immediately obtain full-fledged PKE schemes by using CCA secure SKE [16].

We emphasize that we do not require plaintext awareness and the trapdoor simulatability property to be satisfied by a single building block. This “separation” of the requirements should be contrasted with Dachman-Soled’s construction [18], the building block PKE scheme of which is required to satisfy plaintext awareness and the weak simulatability property simultaneously. We also again emphasize that the assumptions on which both of our constructions are based, are *incomparable* to those used in [18]. Thus, our results add new recipes for constructing CCA secure PKE/KEM from general assumptions (and thus the assumptions that we use could be new targets that are worth pursuing), and also show interesting trade-offs regarding assumptions with Dachman-Soled’s construction.

1.3 Technical Overview

Assumptions on the Building Blocks. *Trapdoor simulatable PKE* (TSPKE) [14] is the key building block for our constructions. TSPKE is a *weaker* (relaxed) version of *simulatable PKE* that was originally formalized by Damgård and Nielsen [19]. Simulatable PKE admits “oblivious sampling” of both public keys and ciphertexts (i.e. sampling them without knowing the randomness or plaintext) in such a way that honestly generated public keys and ciphertexts can be later convincingly explained that they were generated obliviously. These properties are realized by requiring that the key generation algorithm and the encryption algorithm have their own “oblivious sampling” algorithm and its corresponding “inverting” algorithm (where the inverting algorithm corresponds to the algorithm that “explains” that an honest generated public key (or a ciphertext) is sampled obliviously). The difference between TSPKE and simulatable PKE is whether we allow the “inverting” algorithm to take the randomness (and the plaintext) used by the ordinary algorithms (key generation and encryption algorithms) as input. TSPKE allows to take these inputs, while ordinary simulatable PKE does not, which makes the security property of TSPKE weaker but easier to achieve. For our purpose, we only need even a simplified version of TSPKE than the formalization in [14]: we only require a pair (pk, c) of public key/ciphertext (or, a “transcript”) can be obliviously sampled, but not each of pk and c can be so (which is the formalization in [14]). It was shown [19, 14] that we can realize TSPKE from a number of standard cryptographic assumptions, such as the computational and decisional Diffie-Hellman assumptions, RSA, Factoring, and lattice based assumptions. (For more details, see Section 2.2.)

On the other hand, a weakly simulatable PKE scheme (used in the constructions in [43, 18]) considers oblivious sampling only for the encryption algorithm. However, the definition of weakly simulatable PKE used in [43, 18] does not allow the inverting algorithms to take the randomness and the plaintext used by the ordinary encryption algorithm. Therefore, strictly speaking, the “strength” of these primitives as “general cryptographic assumptions” are actually *incomparable*. Nonetheless, the reason why we still think that weakly simulatable PKE

could be viewed as a weaker primitive, is that it does not require the key generation algorithm to be obviously samplable. In fact, this difference is very important for our work. It is this simple difference between TSPKE and weakly simulatable PKE that enables us to weaken the plaintext awareness required in [18], from plaintext awareness in the presence of $O(k)$ keys in [18] into that under only $O(1)$ keys in our constructions.

Ideas for the Constructions. Other than employing TSPKE instead of weakly simulatable PKE, the ideas for our constructions and their security analyses are similar to those in [18]. In particular, the construction of [18] and our constructions are based on the Dolev-Dwork-Naor (DDN) construction [23], but we do not require a non-interactive zero-knowledge proof to ensure the validity of a ciphertext. Instead, the approach of the “double-layered” construction of Myers and Shelat [44] (and its simplifications [31, 37, 40] and variants [38, 39, 41]) is employed, in which a ciphertext consists of the “inner”-layer and “outer”-layer, and the randomness used for generating an outer ciphertext is somehow embedded into an inner ciphertext, so that in the decryption, the validity of the outer ciphertext can be checked by “re-encryption” using the randomness recovered from the inner ciphertext. (In our constructions, the inner-layer encryption is done by a KEM.) In fact, we do a simplification to [18] by removing a one-time signature scheme in [18], by using a commitment scheme, based on the ideas employed in the recent constructions [38, 39, 41].

Recently, Matsuda and Hanaoka [39] introduced the notion of *puncturable tag-based encryption* (PTBE) which abstracts and formalizes the “core” structure of the DDN construction [23]. We define the trapdoor simulatability property for PTBE (and call the primitive *trapdoor simulatable PTBE*) in Section 3, and use this primitive as an “intermediate” building block in our constructions. (This primitive could have other applications than constructing CCA secure PKE, and may be of independent interest.) We also show (in the full version) how to construct a trapdoor simulatable PTBE scheme from a TSPKE scheme. This construction is exactly the same as the construction of a PTBE scheme from a CPA secure PKE scheme used in [39], which is in turn based on the original DDN construction.

Ideas for the Security Proofs. We briefly recall the construction and the security proof in [18], and explain the difference in our proofs and that in [18]. As mentioned above, the construction of [18] is double-layered, where the outer encryption is like the “DDN-lite” construction (i.e. the DDN construction without a non-interactive zero-knowledge proof), and the inner encryption is a multiple-encryption by two PKE schemes. Both the inner and outer encryption schemes use the same building block, with independently generated public keys: $2k$ keys for the outer-layer encryption (that does DDN-lite-encryption) and 2 keys for the inner-layer encryption (that does multiple-encryption by two encryptions). Roughly speaking, in the security proof, [18] constructs a CPA adversary (reduction algorithm) for the inner-layer encryption, from a CCA adversary \mathcal{A} against the entire construction. The reduction algorithm of course has to somehow an-

answer \mathcal{A} 's decryption queries, and this is the place where plaintext awareness comes into play. Plaintext awareness in the ℓ keys setting (sPA1_ℓ security) ensures that for any algorithm \mathcal{C} (called “ciphertext creator”) that receives a set of public keys $(pk_i)_{i \in \{1, \dots, \ell\}}$ and a randomness $r_{\mathcal{C}}$ as input and makes decryption queries, there exists an extractor \mathcal{E} that also receives $(pk_i)_{i \in \{1, \dots, \ell\}}$ and $r_{\mathcal{C}}$ as input, and can “extract” the plaintext from a ciphertext queried by \mathcal{C} . (In our actual security proofs, we denote the “ciphertext creator” by “ \mathcal{A}' ”, but for the explanation here we continue to use \mathcal{C} for clarity.) The idea in the proof in [18] is to use an extractor guaranteed by plaintext awareness to answer the CCA adversary \mathcal{A} 's decryption queries. The problem that arises here is: how do we design the algorithm \mathcal{C} with which the extractor \mathcal{E} is considered? Since the extractor \mathcal{E} needs to be given the randomness $r_{\mathcal{C}}$ used by \mathcal{C} , if we naively design \mathcal{C} , the reduction algorithm cannot use the extractor \mathcal{E} while embedding its instances (the public key and the challenge ciphertext) in the reduction algorithm's CPA security experiment into \mathcal{A} 's view. The approach in [18] is to consider a modified version of the CCA security experiment in which all component ciphertexts (i.e. ciphertexts for the outer-layer encryption) are generated obliviously using some randomness r (which can be performed due to the weak simulatability property of the underlying PKE scheme), and view this modified experiment as a ciphertext creator \mathcal{C} that takes as input $\ell = 2k + 2$ public keys (for both inner-/outer-layer encryptions) and a randomness $r_{\mathcal{C}}$ consisting of the randomness $r_{\mathcal{A}}$ used by \mathcal{A} and the randomness r used for oblivious generation of the component ciphertexts in \mathcal{A} 's challenge ciphertext. ($r_{\mathcal{C}}$ actually also contains some additional randomness used for generating the remaining parts of \mathcal{A} 's challenge ciphertext, but we ignore it here for simplicity.) Designing the algorithm \mathcal{C} in this way, the extractor \mathcal{E} corresponding to \mathcal{C} can be used to answer \mathcal{A} 's decryption queries while the reduction algorithm (attacking the CPA security of the inner-layer encryption) can perform the reduction.

Our main idea for weakening the requirement of plaintext awareness for the building blocks, from $2k + 2$ keys in [18] to $O(1)$ keys, is due to the observation that by relying on the trapdoor simulatability property for the outer-layer encryption, we can “push” the public keys for the outer-layer encryption, into the “randomness” $r_{\mathcal{C}}$ for the ciphertext creator \mathcal{C} (with which the extractor \mathcal{E} is considered), by generating the public keys regarding the outer-layer encryption also obliviously. In order to make this idea work, we thus consider a different design strategy for the ciphertext creator \mathcal{C} . This also enables us to “separate” the requirement that a single building block PKE scheme needs to be simultaneously plaintext aware and simulatable, because we need the simulatability only for the outer-layer encryption.

Actually, like the security proof of the construction in [18], we need to deal with a “bad” decryption query, which is a ciphertext such that its actual decryption result (by the normal decryption algorithm with a secret key) differs from the decryption result obtained by using the extractor \mathcal{E} . (Such a decryption query makes the simulation of the decryption oracle by the reduction algorithm fail.) Our first construction uses the clever trick of Dachman-Soled [18] of using

two CPA secure PKE schemes (that each encrypts a “share” of 2-out-of-2 secret sharing) and their plaintext awareness under 2 keys setting. (As mentioned earlier, in fact, we use a KEM instead of a PKE scheme for the inner encryption.) Dachman-Soled’s approach enables us to use the CPA security and the ability of “detecting” bad queries at the same time. Our second construction is a simplification of our first construction, where we employ a “single” KEM for the inner layer, as opposed to multiple-encryption by two KEMs in our first construction. To detect “bad” decryption queries by an adversary, we employ the ideas and techniques from [44, 31, 37, 40] of using “1-bounded CCA” security [15]. (As mentioned earlier, in fact, CPA security in the presence of one “plaintext-checking” query [47, 1] is sufficient for our purpose.) For more details on these, see Section 4.

1.4 Related Work

The notion of CCA security for PKE was formalized by Naor and Yung [46] and Rackoff and Simon [49]. Since the introduction of the notion, CCA secure PKE schemes have been studied in a number of papers, and thus we only briefly review constructions from general cryptographic assumptions. Dolev, Dwork, and Naor [23] showed the first construction of a CCA secure PKE scheme, from a CPA secure scheme and a NIZK proof system, based on the construction by Naor and Yung [46] that achieves weaker non-adaptive CCA (CCA1) security. These NIZK-based constructions were further improved in [51, 53, 35]. Canetti, Halevi, and Katz [12] showed how to transform an identity-based encryption scheme into a CCA secure PKE scheme. Kiltz [32] showed that the transform of [12] is applicable to a weaker primitive of tag-based encryption (TBE). Peikert and Waters [48] showed how to construct a CCA secure PKE scheme from a *lossy* trapdoor function (TDF). Subsequent works showed that TDFs with weaker security/functionality properties are sufficient for obtaining CCA secure PKE schemes [50, 42, 33, 54, 13]. Hemenway and Ostrovsky [29] showed how to construct a CCA secure scheme in several ways from homomorphic encryption that has some appropriate properties, and the same authors [30] showed that one can construct a CCA secure PKE scheme from a lossy encryption scheme [4] if it can encrypt a plaintext longer than the length of randomness consumed by the encryption algorithm. Myers and Shelat [44] showed that a CCA secure PKE scheme for 1-bit messages can be turned into one with an arbitrarily large plaintext space. Hohenberger, Lewko, and Waters [31] showed that CCA secure PKE can be constructed from a PKE with a weaker security notion called detectable CCA security, from which we can obtain a 1-bit-to-multi-bit transformation for CCA security in a simpler manner than [44]. The simplicity and efficiency of [44] were further improved by Matsuda and Hanaoka [37, 40]. Lin and Tessaro [34] showed how to amplify weak CCA security into strong (ordinary) CCA secure one. Matsuda and Hanaoka [38] showed how to construct a CCA secure PKE scheme by using a CPA secure PKE scheme and point obfuscation [9, 36], and the same authors [39] showed a CCA secure PKE scheme from a CPA secure PKE scheme and a family of hash functions satisfying the

very strong security notion called universal computational extractors (UCE) [3]. The same authors [41] recently also showed that a CCA secure PKE scheme can be built from the combination of a sender non-committing encryption scheme and a key-dependent-message secure SKE scheme. More recently, Hajiabadi and Kapron [28] showed how to construct a CCA secure PKE scheme, from a 1-bit PKE scheme that satisfies circular security and has the structural property called reproducibility.

As has been stated several times, Dachman-Soled [18] showed how to construct a CCA secure PKE scheme from a PKE scheme which simultaneously satisfies weak simulatability [43] and the (standard model) plaintext awareness under the multiple keys setting, which is built based on the result by Myers, Sergi, and Shelat [43] who showed a PKE scheme satisfying the so-called cNM-CCA1 security, from the same building blocks as [18]. Sahai and Waters [52] showed (among other cryptographic primitives) how CCA secure PKE and KEMs can be constructed using an indistinguishability obfuscation [2, 26].

1.5 Paper Organization

In Section 2 (and in Appendix A), we review definitions of primitives and security notions that are necessary for explaining our results. In Section 3, we introduce the notion of trapdoor simulatable PTBE, which is an extension of PTBE introduced in [39], and works as one of main building blocks of our proposed KEMs in the next section. Finally, in Section 4, we show our main results: two constructions of KEMs that show the “trade-off” between “simulatability” property and “plaintext awareness” in Dachman-Soled’s construction [18].

2 Preliminaries

In this section, we review the basic notation and the definitions for plaintext awareness ($\mathsf{sPA}_{1\ell}$ security) [5, 43, 18] of a KEM, trapdoor simulatability properties of a PKE scheme and a commitment scheme, and the syntax of a puncturable tag-based encryption (PTBE) scheme. The definitions for standard cryptographic primitives with standard security definitions that are not reviewed in this section are given in Appendix A, which include PKE, KEMs, and commitment schemes.

Basic Notation. \mathbb{N} denotes the set of all natural numbers, and for $n \in \mathbb{N}$, we define $[n] := \{1, \dots, n\}$. “ $x \leftarrow y$ ” denotes that x is chosen uniformly at random from y if y is a finite set, x is output from y if y is a function or an algorithm, or y is assigned to x otherwise. If x and y are strings, then “ $|x|$ ” denotes the bit-length of x , “ $x||y$ ” denotes the concatenation x and y , and “ $(x \stackrel{?}{=} y)$ ” is the operation which returns 1 if $x = y$ and 0 otherwise. “(P)PTA” stands for a (probabilistic) polynomial time algorithm. For a finite set S , “ $|S|$ ” denotes its size. If \mathcal{A} is a probabilistic algorithm, then “ $y \leftarrow \mathcal{A}(x; r)$ ” denotes that \mathcal{A} computes y as output by taking x as input and using r as randomness, and we just write

“ $y \leftarrow \mathcal{A}(x)$ ” if we do not need to make the randomness used by \mathcal{A} explicit. If furthermore \mathcal{O} is a function or an algorithm, then “ $\mathcal{A}^{\mathcal{O}}$ ” means that \mathcal{A} has oracle access to \mathcal{O} . A function $\epsilon(k) : \mathbb{N} \rightarrow [0, 1]$ is said to be *negligible* if for all positive polynomials $p(k)$ and all sufficiently large $k \in \mathbb{N}$, we have $\epsilon(k) < 1/p(k)$. Throughout this paper, we use the character “ k ” to denote a security parameter.

2.1 Plaintext Awareness for Multiple Keys Setup (sPA1 $_{\ell}$ Security)

Here, we review the definition of (statistical) plaintext awareness for multiple keys setup [43, 18] (denoted by **sPA1 $_{\ell}$ security**, where ℓ denotes the number of keys). Unlike these previous works, we define it for a KEM, rather than a PKE scheme, but we can define plaintext awareness for a KEM in essentially the same way as that for a PKE scheme.

Let $\Gamma = (\text{KKG}, \text{Encap}, \text{Decap})$ be a KEM (where we review the definition of a KEM in Appendix A), and $\ell = \ell(k) > 0$ be a polynomial. Let \mathcal{A} be an algorithm (called a “ciphertext creator”) that takes a set of public keys $(pk_i)_{i \in [\ell]}$ as input, and makes decapsulation queries of the form $(j \in [\ell], c)$ which is supposed to be answered with $K = \text{Decap}(sk_j, c)$. For this \mathcal{A} , we consider the corresponding “(plaintext) extractor” \mathcal{E} : It is a stateful algorithm that initially takes a set of public keys $(pk_i)_{i \in [\ell]}$ and the randomness $r_{\mathcal{A}}$ consumed by \mathcal{A} , and expects to receive “decapsulation” queries of the form $q = (j \in [\ell], c)$; Upon a query, it tries to extract a session-key K corresponding to c so that $K = \text{Decap}(sk_j, c)$, where sk_j is the secret key corresponding to pk_j . After \mathcal{E} extracts a session-key, it may update its internal state to prepare for the next call. Informally, a KEM Γ is said to be **sPA1 $_{\ell}$ secure** if for all PPTA ciphertext creators \mathcal{A} , there exists a corresponding PPTA extractor \mathcal{E} that can work as \mathcal{A} ’s decapsulation oracle in the experiment above.

More formally, for \mathcal{A} that makes $Q = Q(k)$ decapsulation queries, \mathcal{E} , and ℓ , consider the following experiment $\text{Expt}_{\Gamma, \mathcal{A}, \mathcal{E}, \ell}^{\text{sPA1}}(k)$:

$$\begin{aligned} \text{Expt}_{\Gamma, \mathcal{A}, \mathcal{E}, \ell}^{\text{sPA1}}(k) : & [\forall i \in [\ell] : (pk_i, sk_i) \leftarrow \text{KKG}(1^k); r_{\mathcal{A}} \leftarrow \{0, 1\}^*; \\ & \text{st}_{\mathcal{E}} \leftarrow ((pk_i)_{i \in [\ell]}, r_{\mathcal{A}}); \text{Run } \mathcal{A}^{\mathcal{E}(\text{st}_{\mathcal{E}}, \cdot)}((pk_i)_{i \in [\ell]}; r_{\mathcal{A}}) \text{ until it terminates;} \\ & \text{If } \exists i \in [Q] : \text{Decap}(sk_{j_i}, c_i) \neq K_i \text{ then return 1 else return 0.], \end{aligned}$$

where (j_i, c_i) represents \mathcal{A} ’s i -th decapsulation query (which \mathcal{A} expects to be decapsulated as a ciphertext under pk_{j_i}), and K_i represents the answer (i.e. “decapsulation result” of c_i) computed by the algorithm \mathcal{E} . In the experiment, \mathcal{E} is the (possibly stateful) extractor which initially takes $\text{st}_{\mathcal{E}} = ((pk_i)_{i \in [\ell]}, r_{\mathcal{A}})$ as input, and works like \mathcal{A} ’s decapsulation oracle, as explained above.

Definition 1. Let $\ell = \ell(k) > 0$ be a polynomial. We say that a KEM Γ is **sPA1 $_{\ell}$ secure** if for all PPTAs (ciphertext creator) \mathcal{A} , there exists a stateful PPTA (extractor) \mathcal{E} such that $\text{Adv}_{\Gamma, \mathcal{A}, \mathcal{E}, \ell}^{\text{sPA1}}(k) := \Pr[\text{Expt}_{\Gamma, \mathcal{A}, \mathcal{E}, \ell}^{\text{sPA1}}(k) = 1]$ is negligible.

If $\ell = 1$, then **sPA1 $_{\ell}$ security** is equivalent to statistical PA1 security defined by Bellare and Palacio [5]. By definition, trivially, **sPA1 $_x$** implies **sPA1 $_y$** for $x > y$. However, to the best of our knowledge, whether there is an implication (or separation) for the opposite direction, is not known.

2.2 (Simplified) Trapdoor Simulatable Public Key Encryption

Trapdoor simulatable PKE (TSPKE) [14] is a *relaxed* version of simulatable PKE [19]. Simulatable PKE admits “oblivious sampling” of both public keys and ciphertexts (i.e. sampling them without knowing the randomness or plaintext) in such a way that honestly generated public keys and ciphertexts can be later convincingly explained that they were generated obliviously.³ These properties are realized by requiring that the key generation algorithm and the encryption algorithm have their own “oblivious sampling” algorithm and its corresponding “inverting” algorithm (where the inverting algorithm corresponds to the algorithm that explains that an honest generated public key (or a ciphertext) is sampled obviously). The difference between TSPKE and simulatable PKE, is whether we allow for the “inverting” algorithm to take the randomness (and the plaintext) used by the ordinary algorithms PKG and Enc as input. Since the “inverting” algorithm in TSPKE is allowed to see more information than that in simulatable PKE, the former primitive is strictly weaker (and easier to construct) than the latter.

For our purpose, we only need even a simplified version of TSPKE of [14]: we only require a pair (pk, c) of public key/ciphertext (or, “transcript”) can be obliviously sampled [14], but not each of pk and c can be so. A TSPKE scheme with such a simplified syntax may not be useful for constructing non-committing encryption (as done in [19, 14]), but sufficient for our purpose in this paper.

Definition 2. We say that a PKE scheme⁴ $\Pi = (\text{PKG}, \text{Enc}, \text{Dec})$ is trapdoor simulatable (and say that Π is a trapdoor simulatable PKE (TSPKE) scheme) if Π has two additional PPTAs ($\text{oSamp}_\Pi, \text{rSamp}_\Pi$) with the following properties:

- oSamp_Π is the oblivious-sampling algorithm which takes 1^k as input, and outputs an “obliviously generated” public key/ciphertext pair (pk, c) .
- rSamp_Π is the inverting algorithm (corresponding to oSamp_Π) that takes randomness r_g and r_e , and a plaintext m (which are supposed to be used as $(pk, sk) \leftarrow \text{PKG}(1^k; r_g)$ and $c \leftarrow \text{Enc}(pk, m; r_e)$) as input, and outputs a string \hat{r} (that looks like a randomness used by oSamp_Π).
- (**Trapdoor Simulatability**) For all PPTAs $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, $\text{Adv}_{\Pi, \mathcal{A}}^{\text{TSPKE}}(k) := |\Pr[\text{Expt}_{\Pi, \mathcal{A}}^{\text{TSPKE-Real}}(k) = 1] - \Pr[\text{Expt}_{\Pi, \mathcal{A}}^{\text{TSPKE-Sim}}(k) = 1]|$ is negligible, where the experiments $\text{Expt}_{\Pi, \mathcal{A}}^{\text{TSPKE-Real}}(k)$ and $\text{Expt}_{\Pi, \mathcal{A}}^{\text{TSPKE-Sim}}(k)$ are defined as in Fig. 1 (upper-left and upper-right, respectively).

Concrete Instantiations of TSPKE. Since our definition of TSPKE is a simplified (and hence weaker) version of the definition by Choi et al. [14], and TSPKE is a weaker primitive than a simulatable PKE scheme in the sense of Damgård and Nielsen [19], we can use any of (trapdoor) simulatable PKE schemes shown in these works. In particular, we can construct a TSPKE scheme from most of the

³ (Trapdoor) simulatable PKE scheme was introduced as a building block for constructing non-committing encryption [11].

⁴ The syntax of PKE is reviewed in Appendix A.

$\text{Expt}_{\Pi, \mathcal{A}}^{\text{TSPKE-Real}}(k) :$ $(m, \text{st}) \leftarrow \mathcal{A}_1(1^k)$ $r_g, r_e \leftarrow \{0, 1\}^*$ $(pk, sk) \leftarrow \text{PKG}(1^k; r_g)$ $c \leftarrow \text{Enc}(pk, m; r_e)$ $\hat{r} \leftarrow \text{rSamp}_{\Pi}(r_g, r_e, m)$ $b' \leftarrow \mathcal{A}_2(\text{st}, pk, c, \hat{r})$ $\text{Return } b'.$	$\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{TSPKE-Sim}}(k) :$ $(m, \text{st}) \leftarrow \mathcal{A}_1(1^k)$ $\hat{r} \leftarrow \{0, 1\}^*$ $(pk, c) \leftarrow \text{oSamp}_{\Pi}(1^k; \hat{r})$ $b' \leftarrow \mathcal{A}_2(\text{st}, pk, c, \hat{r})$ $\text{Return } b'.$
$\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{TSPBE-Real}}(k) :$ $(\text{tag}^*, m, \text{st}) \leftarrow \mathcal{A}_1(1^k)$ $r_g, r_e \leftarrow \{0, 1\}^*$ $(pk, sk) \leftarrow \text{TKG}(1^k; r_g)$ $c \leftarrow \text{TEnc}(pk, \text{tag}^*, m; r_e)$ $\widehat{sk}_{\text{tag}^*} \leftarrow \text{Punc}(sk, \text{tag}^*)$ $\hat{r} \leftarrow \text{rSamp}_{\mathcal{T}}(r_g, r_e, \text{tag}^*, m)$ $b' \leftarrow \mathcal{A}_2(\text{st}, pk, c, \widehat{sk}_{\text{tag}^*}, \hat{r})$ $\text{Return } b'.$	$\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{TSPBE-Sim}}(k) :$ $(\text{tag}^*, m, \text{st}) \leftarrow \mathcal{A}_1(1^k)$ $\hat{r} \leftarrow \{0, 1\}^*$ $(pk, c, \widehat{sk}_{\text{tag}^*}) \leftarrow \text{oSamp}_{\mathcal{T}}(\text{tag}^*; \hat{r})$ $b' \leftarrow \mathcal{A}_2(\text{st}, pk, c, \widehat{sk}_{\text{tag}^*}, \hat{r})$ $\text{Return } b'.$

Fig. 1. Security experiments for defining security of TSPKE (upper-left and upper-right) and those for defining security of TSPTBE (bottom-left and bottom-right)

standard cryptographic assumptions such as the computational and decisional Diffie-Hellman, RSA, factoring, and learning-with-errors assumptions [19, 14]. (For example, the ElGamal encryption, Damgård’s ElGamal encryption, and Cramer-Shoup-Lite encryption schemes can be shown to be a TSPKE scheme if they are implemented in a simulatable group [20].) In terms of “general” cryptographic assumptions, Damgård and Nielsen [19] showed that a simulatable PKE scheme can be constructed from a family of trapdoor permutations with the simulatability property, in which the key generation and the domain-sampling algorithms have the oblivious sampling property (which is defined analogously to simulatable PKE). Hence, we can also construct a TSPKE from it.

2.3 Trapdoor Simulatable Commitment Schemes

Let $\mathcal{C} = (\text{CKG}, \text{Com})$ be a commitment scheme. (We review the syntax of a commitment scheme and its “target-binding” property in Appendix A.)

We define the trapdoor simulatability property of a commitment scheme \mathcal{C} in exactly the same way as the trapdoor simulatability of a PKE scheme. Namely, we require that there be the oblivious sampling algorithm $\text{oSamp}_{\mathcal{C}}$ (for sampling a key/commitment pair (ck, c)) and the corresponding inverting algorithm $\text{rSamp}_{\mathcal{C}}$, whose interfaces are exactly the same as oSamp_{Π} and rSamp_{Π} of a TSPKE scheme, respectively. We say that a commitment scheme \mathcal{C} is *trapdoor simulatable* (and say that \mathcal{C} is a trapdoor simulatable commitment scheme) if for all PPTA adversaries \mathcal{A} , the advantage $\text{Adv}_{\mathcal{C}, \mathcal{A}}^{\text{TSCom}}(k) := |\Pr[\text{Expt}_{\mathcal{C}, \mathcal{A}}^{\text{TSCom-Real}}(k) = 1] - \Pr[\text{Expt}_{\mathcal{C}, \mathcal{A}}^{\text{TSCom-Sim}}(k) = 1]|$ is negligible, where the experiments $\text{Expt}_{\mathcal{C}, \mathcal{A}}^{\text{TSCom-Real}}(k)$ and $\text{Expt}_{\mathcal{C}, \mathcal{A}}^{\text{TSCom-Sim}}(k)$ are defined in exactly the same way as $\text{Expt}_{\Pi, \mathcal{A}}^{\text{TSPKE-Real}}(k)$ and

$\text{Expt}_{\mathcal{C}, \mathcal{A}}^{\text{TSPKE-Sim}}(k)$ for a TSPKE scheme, respectively (and thus we do not write down them).

We can achieve a commitment scheme which satisfies target-binding, trapdoor simulatability, and the requirement of the size of commitments (namely we require the size of commitments to be k -bit for k -bit security), only from a TSPKE scheme and a universal one-way hash function (UOWHF) [45], just by hashing a ciphertext of the TSPKE scheme by the UOWHF. This construction is given in the full version.

2.4 Puncturable Tag-Based Encryption

Here, we recall the syntax of puncturable tag-based encryption (PTBE), which was introduced by Matsuda and Hanaoka [39] as an abstraction of the “core” structure of the Dolev-Dwork-Naor (DDN) construction [23]. Similarly to [39], we use PTBE as an intermediate building block to reduce the description complexity of our proposed constructions in Section 4.

Intuitively, a PTBE scheme is a TBE scheme that has a mechanism for generating a “punctured” secret key $\widehat{sk}_{\text{tag}^*}$, according to a “punctured point” tag tag^* . The punctured secret key can be used to decrypt all “honestly generated” ciphertexts that are generated under tags that are different from tag^* , while the punctured secret key is useless for decrypting ciphertexts generated under tag^* .

Formally, a PTBE scheme consists of the five PPTAs (TKG, TEnc, TDec, Punc, $\widehat{\text{TDec}}$) among which the latter three algorithms are deterministic, with the following interface:

$$\begin{array}{lll} \textbf{Key Generation:} & \textbf{Encryption:} & \textbf{Decryption:} \\ (pk, sk) \leftarrow \text{TKG}(1^k) & c \leftarrow \text{TEnc}(pk, \text{tag}, m) & m \text{ (or } \perp) \leftarrow \text{TDec}(sk, \text{tag}, c) \\ \\ \textbf{Puncturing:} & & \textbf{Punctured Decryption:} \\ \widehat{sk}_{\text{tag}^*} \leftarrow \text{Punc}(sk, \text{tag}^*) & & m \text{ (or } \perp) \leftarrow \widehat{\text{TDec}}(\widehat{sk}_{\text{tag}^*}, \text{tag}, c) \end{array}$$

where (pk, sk) is a public/secret key pair, c is a ciphertext of a plaintext m under pk and a tag $\text{tag} \in \{0, 1\}^k$, and $\widehat{sk}_{\text{tag}^*}$ is a “punctured” secret key corresponding to a tag $\text{tag}^* \in \{0, 1\}^k$.

We require for all $k \in \mathbb{N}$, all tags $\text{tag}^*, \text{tag} \in \{0, 1\}^k$ such that $\text{tag}^* \neq \text{tag}$, all (pk, sk) output from $\text{TKG}(1^k)$, all plaintexts m , and all ciphertexts c output from $\text{TEnc}(pk, \text{tag}, m)$, it holds that $\text{TDec}(sk, \text{tag}, c) = \widehat{\text{TDec}}(\widehat{\text{Punc}}(sk, \text{tag}^*), \text{tag}, c) = m$.

In [39], the security notion called “extended CPA security” was defined as a security notion of PTBE. In our proposed KEMs, we need a stronger security property for PTBE, which is an analogue of TSPKE, and we will introduce it in the next section.

3 Trapdoor Simulatable PTBE

In this section, we define trapdoor simulatability of a PTBE scheme, in the same way as that of a PKE scheme and a commitment scheme. However, for

the oblivious sampling algorithm, we let it take a “punctured point” tag tag^* as input, and require that it output the punctured secret key $\widehat{sk}_{\text{tag}^*}$ (corresponding to tag^*) in addition to a public key/ciphertext pair (pk, c) .

Formally, we define a trapdoor simulatable PTBE (TSPTBE) as follows:

Definition 3. We say that a PTBE scheme $\mathcal{T} = (\text{TKG}, \text{TEnc}, \text{TDec}, \text{Punc}, \widehat{\text{TDec}})$ is trapdoor simulatable (and say that \mathcal{T} is a trapdoor simulatable PTBE (TSPTBE) scheme) if \mathcal{T} has two additional PPTAs ($\text{oSamp}_{\mathcal{T}}, \text{rSamp}_{\mathcal{T}}$) with the following properties:

- $\text{oSamp}_{\mathcal{T}}$ is the oblivious sampling algorithm which takes a “punctured point” tag tag^* as input, and outputs an “obliviously generated” public key/ciphertext pair (pk, c) and a punctured secret key $\widehat{sk}_{\text{tag}^*}$.
- $\text{rSamp}_{\mathcal{T}}$ is the inverting algorithm (corresponding to $\text{oSamp}_{\mathcal{T}}$) that takes 1^k , randomness r_g and r_e , a “punctured point” tag tag^* , and a plaintext m (which are supposed to be used as $(pk, sk) \leftarrow \text{TKG}(1^k; r_g)$ and $c \leftarrow \text{TEnc}(pk, \text{tag}^*, m; r_e)$) as input, and outputs a string \widehat{r} (that looks like a randomness used by $\text{oSamp}_{\mathcal{T}}$).
- (**Trapdoor Simulatability**) For all PPTAs $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, $\text{Adv}_{\mathcal{T}, \mathcal{A}}^{\text{TSPTBE}}(k) := |\Pr[\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{TSPTBE-Real}}(k) = 1] - \Pr[\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{TSPTBE-Sim}}(k) = 1]|$ is negligible, where the experiments $\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{TSPTBE-Real}}(k)$ and $\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{TSPTBE-Sim}}(k)$ are defined as in Fig. 1 (bottom-left and bottom-right, respectively).

On the Existence of TSPTBE. Though it might look complicated, we can construct a TSPTBE scheme from a TSPKE scheme, by a Dolev-Dwork-Naor-style approach [23]. The construction is exactly the same as the construction of a PTBE scheme from any CPA secure PKE shown in [39], which is the “core” structure of the DDN construction, namely, the DDN construction without a NIZK proof and without its one-time signature. (For this construction, we can straightforwardly consider the oblivious sampling algorithm and the corresponding inverting algorithm.) We prove the following lemma in the full version.

Lemma 1. *If a TSPKE scheme exists, then so does a TSPTBE scheme.*

Useful fact. For the security proofs of our constructions in Section 4, we will use the fact that the straightforward concatenation of a “transcript” of a trapdoor simulatable commitment and that of a TSPTBE scheme, also admits the trapdoor simulatable property.

More formally, for a TSPTBE scheme $\mathcal{T} = (\text{TKG}, \text{TEnc}, \text{TDec}, \text{Punc}, \widehat{\text{TDec}}, \text{oSamp}_{\mathcal{T}}, \text{rSamp}_{\mathcal{T}})$ and a trapdoor simulatable commitment scheme $\mathcal{C} = (\text{CKG}, \text{Com}, \text{oSamp}_{\mathcal{C}}, \text{rSamp}_{\mathcal{C}})$ such that the plaintext space of \mathcal{T} and that of \mathcal{C} are identical, and for an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, consider the following “real” experiment $\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{A}}^{\text{TS-Real}}(k)$ and the “simulated” experiment $\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{A}}^{\text{TS-Sim}}(k)$ as described in Fig. 2 (left and right, respectively).

Then, we can prove the following lemma, whose proof is almost straightforward due to the trapdoor simulatability property of \mathcal{C} and \mathcal{T} . The proof is by a standard hybrid argument, and is given in the full version.

$\text{Expt}_{[C, \mathcal{T}], \mathcal{A}}^{\text{TS-Real}}(k) :$ $(m, \text{st}) \leftarrow \mathcal{A}_1(1^k)$ $r_g, r'_g, r_c, r_t \leftarrow \{0, 1\}^*$ $ck \leftarrow \text{CKG}(1^k; r_g)$ $\text{tag}^* \leftarrow \text{Com}(ck, m; r_c)$ $\widehat{r}_c \leftarrow \text{rSamp}_{\mathcal{C}}(r_g, r_c, m)$ $(pk, sk) \leftarrow \text{TKG}(1^k; r'_g)$ $c^* \leftarrow \text{TEnc}(ck, \text{tag}^*, m; r_e)$ $\widehat{sk}_{\text{tag}^*} \leftarrow \text{Punc}(sk, \text{tag}^*)$ $\widehat{r}_t \leftarrow \text{rSamp}_{\mathcal{T}}(r'_g, r_t, \text{tag}^*, m)$ $b' \leftarrow \mathcal{A}_2(\text{st}, ck, \text{tag}^*, pk, c^*, \widehat{sk}_{\text{tag}^*}, \widehat{r}_c, \widehat{r}_t)$ $\text{Return } b'.$	$\text{Expt}_{[C, \mathcal{T}], \mathcal{A}}^{\text{TS-Sim}}(k) :$ $(m, \text{st}) \leftarrow \mathcal{A}_1(1^k)$ $\widehat{r}_c, \widehat{r}_t \leftarrow \{0, 1\}^*$ $(ck, \text{tag}^*) \leftarrow \text{oSamp}_{\mathcal{C}}(1^k; \widehat{r}_c)$ $(pk, c^*, \widehat{sk}_{\text{tag}^*}) \leftarrow \text{oSamp}_{\mathcal{T}}(\text{tag}^*; \widehat{r}_t)$ $b' \leftarrow \mathcal{A}_2(\text{st}, ck, \text{tag}^*, pk, c^*, \widehat{sk}_{\text{tag}^*}, \widehat{r}_c, \widehat{r}_t)$ $\text{Return } b'.$
---	---

Fig. 2. Security experiments for defining the trapdoor simulatability of the concatenation of a “transcript” of a commitment scheme and that of a TSPTBE scheme.

Lemma 2. *Assume that the commitment scheme \mathcal{C} and the PTBE scheme \mathcal{T} are trapdoor simulatable. Then, for all PPTAs \mathcal{A} , $\text{Adv}_{[C, \mathcal{T}], \mathcal{A}}^{\text{TS}}(k) := |\Pr[\text{Expt}_{[C, \mathcal{T}], \mathcal{A}}^{\text{TS-Real}}(k) = 1] - \Pr[\text{Expt}_{[C, \mathcal{T}], \mathcal{A}}^{\text{TS-Sim}}(k) = 1]|$ is negligible.*

4 Proposed KEMs

In this section, we show our main results: two KEMs that show the “trade-off” between the strength of (standard model) plaintext awareness and the simulatability property with those of the construction by Dachman-Soled [18].

In Section 4.1, we show our first construction, which is CCA secure based on a KEM satisfying CPA security and sPA1₂ security, and a TSPKE scheme. In Section 4.2, we show our second construction which is CCA secure based on a KEM satisfying 1-CCA security and sPA1₁ security, and a TSPKE scheme.

4.1 First Construction

Let $\Gamma_{\text{in}} = (\text{KKG}_{\text{in}}, \text{Encap}_{\text{in}}, \text{Decap}_{\text{in}})$ be a KEM whose ciphertext length is $n = n(k)$ and whose session-key space is $\{0, 1\}^{3k}$ for k -bit security.⁵ Let $\mathcal{T} = (\text{TKG}, \text{TEnc}, \text{TDec}, \text{Punc}, \widehat{\text{TDec}})$ be a PTBE scheme and $\mathcal{C} = (\text{CKG}, \text{Com})$ be a commitment scheme. We require the plaintext space of TEnc and the message space of Com to be $\{0, 1\}^{2n}$, and the randomness space of TEnc and that of Com to be $\{0, 1\}^k$ for k -bit security.⁶ Then, our first proposed KEM $\Gamma = (\text{KKG}, \text{Encap}, \text{Decap})$ is constructed as in Fig. 3.

⁵ Note that the session-key space of a KEM can be adjusted “for free” by applying a pseudorandom generator to a session-key. Such a construction preserves CPA and sPA1_ℓ security.

⁶ The requirements of the randomness space of TEnc and Com are without loss of generality, because we can adjust them using a pseudorandom generator. (The trapdoor simulatability property is preserved even if we use a pseudorandom generator.)

$\text{KKG}(1^k) :$ $(pk_{\text{in}0}, sk_{\text{in}0}) \leftarrow \text{KKG}_{\text{in}}(1^k)$ $(pk_{\text{in}1}, sk_{\text{in}1}) \leftarrow \text{KKG}_{\text{in}}(1^k)$ $(pk, sk) \leftarrow \text{TKG}(1^k)$ $ck \leftarrow \text{CKG}(1^k)$ $PK \leftarrow (pk_{\text{in}0}, pk_{\text{in}1}, pk, ck)$ $SK \leftarrow (sk_{\text{in}0}, sk_{\text{in}1}, sk, PK)$ Return (PK, SK) .	$\text{Decap}(SK, C) :$ $(sk_{\text{in}0}, sk_{\text{in}1}, sk, PK) \leftarrow SK$ $(pk_{\text{in}0}, pk_{\text{in}1}, pk, ck) \leftarrow PK$ $(\text{tag}, c) \leftarrow C$ $(c_{\text{in}0} \ c_{\text{in}1}) \leftarrow \text{TDec}(sk, \text{tag}, c)$ If TDec has returned \perp then return \perp . $\alpha_0 \leftarrow \text{Decap}_{\text{in}}(sk_{\text{in}0}, c_{\text{in}0})$ $\alpha_1 \leftarrow \text{Decap}_{\text{in}}(sk_{\text{in}1}, c_{\text{in}1})$ If $\alpha_0 = \perp$ or $\alpha_1 = \perp$ then return \perp . $\alpha \leftarrow \alpha_0 \oplus \alpha_1$ Parse α as $(r_c, r_t, K) \in (\{0, 1\}^k)^3$ If $\text{Com}(ck, (c_{\text{in}0} \ c_{\text{in}1}); r_c) = \text{tag}$ and $\text{TEnc}(pk, \text{tag}, (c_{\text{in}0} \ c_{\text{in}1}); r_t) = c$ then return K else return \perp .
$\text{Encap}(PK) :$ $(pk_{\text{in}0}, pk_{\text{in}1}, pk, ck) \leftarrow PK$ $(c_{\text{in}0}, \alpha_0) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}0})$ $(c_{\text{in}1}, \alpha_1) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}1})$ $\alpha \leftarrow \alpha_0 \oplus \alpha_1$ Parse α as $(r_c, r_t, K) \in (\{0, 1\}^k)^3$ $\text{tag} \leftarrow \text{Com}(ck, (c_{\text{in}0} \ c_{\text{in}1}); r_c)$ $c \leftarrow \text{TEnc}(pk, \text{tag}, (c_{\text{in}0} \ c_{\text{in}1}); r_t)$ $C \leftarrow (\text{tag}, c)$. Return (C, K) .	

Fig. 3. The first proposed construction: the KEM Γ based on a KEM Γ_{in} , a commitment scheme \mathcal{C} , and a PTBE scheme \mathcal{T} .

Alternative Decapsulation Algorithm. Similarly to the constructions in [37–39], to show the CCA security of the proposed KEM Γ , it is useful to consider the following alternative decapsulation algorithm AltDecap . For a k -bit string $\text{tag}^* \in \{0, 1\}^k$ and a key pair (PK, SK) output by $\text{KKG}(1^k)$, where $PK = (pk_{\text{in}0}, pk_{\text{in}1}, pk, ck)$ and $SK = (sk_{\text{in}0}, sk_{\text{in}1}, sk, PK)$, we define an “alternative” secret key $\widehat{SK}_{\text{tag}^*}$ associated with $\text{tag}^* \in \{0, 1\}^k$ by $\widehat{SK}_{\text{tag}^*} = (sk_{\text{in}0}, sk_{\text{in}1}, \text{tag}^*, \widehat{sk}_{\text{tag}^*}, PK)$, where $\widehat{sk}_{\text{tag}^*} = \text{Punc}(sk, \text{tag}^*)$. AltDecap takes an “alternative” secret key $\widehat{SK}_{\text{tag}^*}$ defined as above and a ciphertext $C = (\text{tag}, c)$ as input, and runs as follows:

$\text{AltDecap}(\widehat{SK}_{\text{tag}^*}, C)$: First check if $\text{tag}^* = \text{tag}$, and return \perp if this is the case. Otherwise, run in exactly the same way as $\text{Decap}(SK, C)$, except that “ $(c_{\text{in}0} \| c_{\text{in}1}) \leftarrow \widehat{\text{TDec}}(\widehat{sk}_{\text{tag}^*}, \text{tag}, c)$ ” is executed in the fourth step, instead of “ $(c_{\text{in}0} \| c_{\text{in}1}) \leftarrow \text{TDec}(sk, \text{tag}, c)$.”

Regarding AltDecap , the following lemma is easy to see due to the correctness of the underlying PTBE scheme \mathcal{T} and the validity check of c by re-encryption performed at the last step. (The formal proof is given in the full version.)

Lemma 3. *Let $\text{tag}^* \in \{0, 1\}^k$ be a string and let (PK, SK) be a key pair output by $\text{KKG}(1^k)$. Furthermore, let $\widehat{SK}_{\text{tag}^*}$ be an alternative secret key as defined above. Then, for any ciphertext $C = (\text{tag}, c)$ (which could be outside the range of $\text{Encap}(PK)$) satisfying $\text{tag} \neq \text{tag}^*$, it holds that $\text{Decap}(SK, C) = \text{AltDecap}(\widehat{SK}_{\text{tag}^*}, C)$.*

CCA Security. The security of Γ is guaranteed by the following theorem.

Theorem 1. *Assume that the KEM Γ_{in} is CPA secure and sPA1₂ secure, the commitment scheme \mathcal{C} is target-binding and trapdoor simulatable, and the PTBE scheme \mathcal{T} is trapdoor simulatable. Then, the KEM Γ constructed as in Fig. 3 is CCA secure.*

Note that as mentioned in Section 2.3, a commitment scheme with trapdoor simulatability and target-binding can be constructed from any TSPKE scheme, and thus the above theorem shows that we can indeed construct a CCA secure KEM (and thus CCA secure PKE) from the combination of a KEM satisfying CPA and sPA1₂ security and a TSPKE scheme.

We have provided ideas for the security proof in Section 1.3, and thus we directly proceed to the proof.

Proof of Theorem 1. Let \mathcal{A} be any PPTA adversary that attacks the CCA security of the KEM Γ . Our security proof is via the sequence of games argument. To describe the games, we will need an extractor \mathcal{E} corresponding to some “ciphertext creator” \mathcal{A}' that is guaranteed to exist by the sPA1₂ security of Γ_{in} . Specifically, consider the following \mathcal{A}' (that internally runs \mathcal{A}) that runs in the experiment $\text{Expt}_{\Gamma_{\text{in}}, \mathcal{A}', \mathcal{E}, 2}^{\text{sPA1}}$ (k), with a corresponding extractor \mathcal{E} :

$\mathcal{A}'^{\mathcal{E}(\text{st}_{\mathcal{E}}, \cdot)}(pk_1, pk_2; r_{\mathcal{A}'} = (r_{\mathcal{A}}, \widehat{r}_c, \widehat{r}_t, K^*))$: \mathcal{A}' firstly sets $pk_{\text{in}0} \leftarrow pk_1$ and $pk_{\text{in}1} \leftarrow pk_2$ (which implicitly sets $sk_{\text{in}0} \leftarrow sk_1$ and $sk_{\text{in}1} \leftarrow sk_2$, where sk_1 (resp. sk_2) is the secret key corresponding to pk_1 (resp. pk_2)), and runs $(ck, \text{tag}^*) \leftarrow \text{oSamp}_{\mathcal{C}}(1^k; \widehat{r}_c)$ and $(pk, c^*, \widehat{sk}_{\text{tag}^*}) \leftarrow \text{oSamp}_{\mathcal{T}}(\text{tag}^*; \widehat{r}_t)$. Then \mathcal{A}' sets $PK \leftarrow (pk_{\text{in}0}, pk_{\text{in}1}, pk, ck)$ and $C^* \leftarrow (\text{tag}^*, c^*)$, and then runs $\mathcal{A}(PK, C^*, K^*; r_{\mathcal{A}})$. When \mathcal{A} submits a decapsulation query C , \mathcal{A}' responds to it as if it runs $\text{AltDecap}(\widehat{SK}_{\text{tag}^*}, C)$, where the oracle calls (to the extractor \mathcal{E}) of the form $(1, c_{\text{in}0})$ and $(2, c_{\text{in}1})$ are used as substitutes for $\text{Decap}_{\text{in}}(sk_{\text{in}0}, c_{\text{in}0})$ and $\text{Decap}_{\text{in}}(sk_{\text{in}1}, c_{\text{in}1})$, respectively. More precisely, \mathcal{A}' answers \mathcal{A} 's decapsulation query $C = (\text{tag}, c)$ as follows:

1. If $\text{tag} = \text{tag}^*$, then return \perp to \mathcal{A} .
2. Run $(c_{\text{in}0} \| c_{\text{in}1}) \leftarrow \widehat{\text{TDec}}(\widehat{sk}_{\text{tag}^*}, \text{tag}, c)$, and return \perp to \mathcal{A} if $\widehat{\text{TDec}}$ has returned \perp .
3. Submit queries $(1, c_{\text{in}0})$ and $(2, c_{\text{in}1})$ to the extractor $\mathcal{E}(\text{st}_{\mathcal{E}}, \cdot)$ and receive the answers α_0 and α_1 , respectively. (Here, the answers α_0 and α_1 are expected to be $\alpha_0 = \text{Decap}_{\text{in}}(sk_{\text{in}0}, c_{\text{in}0})$ and $\alpha_1 = \text{Decap}_{\text{in}}(sk_{\text{in}1}, c_{\text{in}1})$, respectively, and the extractor \mathcal{E} may update its state upon each call.)
4. If $\alpha_0 = \perp$ or $\alpha_1 = \perp$, then return \perp to \mathcal{A} .
5. Let $\alpha \leftarrow \alpha_0 \oplus \alpha_1$ and parse α as $(r_c, r_t, K) \in (\{0, 1\}^k)^3$.
6. If $\text{Com}(ck, (c_{\text{in}0} \| c_{\text{in}1}); r_c) = \text{tag}$ and $\text{TEnc}(pk, (c_{\text{in}0} \| c_{\text{in}1}); r_t) = c$, then return K , otherwise return \perp , to \mathcal{A} .

When \mathcal{A} terminates, \mathcal{A}' also terminates.

The above completes the description of the algorithm \mathcal{A}' . The randomness $r_{\mathcal{A}'}$ consumed by \mathcal{A}' is of the form $(r_{\mathcal{A}}, \widehat{r}_c, \widehat{r}_t, K^*)$, where $r_{\mathcal{A}}$, \widehat{r}_c , and \widehat{r}_t are the randomness used by \mathcal{A} , $\text{oSamp}_{\mathcal{C}}$, and $\text{oSamp}_{\mathcal{T}}$, respectively, and K^* is a k -bit

string. The corresponding extractor \mathcal{E} thus receives (pk_1, pk_2) and $r_{\mathcal{A}'}$ as its initial state $\text{st}_{\mathcal{E}}$. Note that since Γ_{in} is assumed to be sPA1_2 secure and \mathcal{A}' is a PPTA, $\text{Adv}_{\Gamma_{\text{in}}, \mathcal{A}', \mathcal{E}, 2}^{\text{sPA1}}(k)$ is negligible for this extractor \mathcal{E} , which will be used later in the proof. (Looking ahead, we will design the sequence of games so that \mathcal{A} 's view in the case \mathcal{A} is internally run by \mathcal{A}' and \mathcal{A}' is run in $\text{Expt}_{\Gamma_{\text{in}}, \mathcal{A}', \mathcal{E}, 2}^{\text{sPA1}}(k)$, is identical to \mathcal{A} 's view in Game 6.)

For convenience, we refer to the procedure of using the extractor \mathcal{E} as substitutes for $\text{Decap}_{\text{in}}(sk_{\text{in}0}, \cdot)$ and $\text{Decap}_{\text{in}}(sk_{\text{in}1}, \cdot)$, as $\text{AltDecap}'_{\mathcal{E}}$. Here, $\text{AltDecap}'_{\mathcal{E}}$ is a stateful procedure that initially takes tag^* , $\widehat{sk}_{\text{tag}^*}$, and an initial state $\text{st}_{\mathcal{E}}$ of \mathcal{E} (i.e. $\text{st}_{\mathcal{E}} = ((pk_{\text{in}0}, pk_{\text{in}1}), r_{\mathcal{A}'})$) as input, and expects to receive a ciphertext $C = (\text{tag}, c)$ as an input. If it receives a ciphertext $C = (\text{tag}, c)$, it calculates the decapsulation result K (or \perp) as \mathcal{A}' does for \mathcal{A} , using $\widehat{sk}_{\text{tag}^*}$ and the extractor \mathcal{E} , where \mathcal{E} 's internal state could be updated upon each execution.

Now, using the adversary \mathcal{A} and the extractor \mathcal{E} , consider the following sequence of games: (Here, the values with asterisk (*) represent those related to the challenge ciphertext for \mathcal{A} .)

- Game 1:** This is the experiment $\text{Expt}_{\Gamma, \mathcal{A}}^{\text{CCA}}(k)$ itself.
- Game 2:** Same as Game 1, except that all decapsulation queries $C = (\text{tag}, c)$ satisfying $\text{tag} = \text{tag}^*$ are answered with \perp .
- Game 3:** Same as Game 2, except that all decapsulation queries C are answered with $\text{AltDecap}(\widehat{SK}_{\text{tag}^*}, C)$, where $\widehat{SK}_{\text{tag}^*}$ is the alternative secret key corresponding to (PK, SK) and tag^* . Furthermore, we pick a random bit $\gamma \in \{0, 1\}$ uniformly at random just before executing \mathcal{A} , which will be used to define the events in this game and the subsequent games. (γ does not appear in \mathcal{A} 's view in this and all subsequent games, and thus does not affect its behavior at all.)
- Game 4:** In this game, we use $\text{AltDecap}'_{\mathcal{E}}$ (defined as above) as \mathcal{A} 's decapsulation oracle, where the initial state of \mathcal{E} (used internally by $\text{AltDecap}'_{\mathcal{E}}$) is prepared using the “inverting algorithms” $\text{rSamp}_{\mathcal{C}}$ of \mathcal{C} and $\text{rSamp}_{\mathcal{T}}$ of \mathcal{T} . Moreover, we also change the ordering of the steps so that they do not affect \mathcal{A} 's view.

More precisely, this game is defined as follows:

Game 4:	$\begin{aligned} & (pk_{\text{in}0}, sk_{\text{in}0}) \leftarrow \text{KKG}_{\text{in}}(1^k); \\ & (pk_{\text{in}1}, sk_{\text{in}1}) \leftarrow \text{KKG}_{\text{in}}(1^k); \\ & (c_{\text{in}0}^*, \alpha_0^*) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}0}); \\ & (c_{\text{in}1}^*, \alpha_1^*) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}1}); \\ & \alpha^* \leftarrow (\alpha_0^* \oplus \alpha_1^*); \\ & \text{Parse } \alpha^* \text{ as } (r_c^*, r_t^*, K_1^*) \in (\{0, 1\}^k)^3; \\ & r_g \leftarrow \{0, 1\}^*; \\ & ck \leftarrow \text{CKG}(1^k; r_g); \\ & \text{tag}^* \leftarrow \text{Com}(ck, (c_{\text{in}0}^* \ c_{\text{in}1}^*); r_c^*); \\ & \widehat{r}_c \leftarrow \text{rSamp}_{\mathcal{C}}(r_g, r_c^*, (c_{\text{in}0}^* \ c_{\text{in}1}^*)); \\ & (\text{Continue to the right column } \nearrow) \end{aligned}$	$\begin{aligned} & r'_g \leftarrow \{0, 1\}^*; \\ & (pk, sk) \leftarrow \text{TKG}(1^k; r'_g); \\ & \widehat{sk}_{\text{tag}^*} \leftarrow \text{Punc}(sk, \text{tag}^*); \\ & c^* \leftarrow \text{TEnc}(pk, \text{tag}^*, (c_{\text{in}0}^* \ c_{\text{in}1}^*); r_t^*); \\ & \widehat{r}_t \leftarrow \text{rSamp}_{\mathcal{T}}(r'_g, r_t^*, \text{tag}^*, (c_{\text{in}0}^* \ c_{\text{in}1}^*)); \\ & PK \leftarrow (pk_{\text{in}0}, pk_{\text{in}1}, pk, ck); \\ & C^* \leftarrow (\text{tag}^*, c^*); \\ & K_0^* \leftarrow \{0, 1\}^k; \\ & b \leftarrow \{0, 1\}; \\ & r_{\mathcal{A}} \leftarrow \{0, 1\}^*; \\ & r_{\mathcal{A}'} \leftarrow (r_{\mathcal{A}}, \widehat{r}_c, \widehat{r}_t, K_b^*); \\ & \text{st}_{\mathcal{E}} \leftarrow ((pk_{\text{in}0}, pk_{\text{in}1}), r_{\mathcal{A}'}); \\ & \gamma \leftarrow \{0, 1\}; \\ & b' \leftarrow \mathcal{A}^{\mathcal{O}}(PK, C^*, K_b^*; r_{\mathcal{A}}) \end{aligned}$
----------------	---	---

where the decapsulation oracle \mathcal{O} that \mathcal{A} has access in Game 4 is $\text{AltDecap}'_{\mathcal{E}}$ (which initially receives $\text{tag}^*, \widehat{sk}_{\text{tag}^*}, \text{st}_{\mathcal{E}} = (pk_{\text{in}0}, pk_{\text{in}1}, r_{\mathcal{A}'})$ as input). Note that the extractor \mathcal{E} used internally by $\text{AltDecap}'_{\mathcal{E}}$ may update its state $\text{st}_{\mathcal{E}}$ upon each execution.

Game 5: Same as Game 4, except that $r_c^*, r_t^*, K_1^* \in \{0, 1\}^k$ are picked uniformly at random, independently of $\alpha^* = \alpha_0^* \oplus \alpha_1^*$. That is, the steps “ $\alpha^* \leftarrow \alpha_0^* \oplus \alpha_1^*$; Parse α^* as $(r_c^*, r_t^*, K_1^*) \in (\{0, 1\}^k)^3$ ” in Game 4 are replaced with the step “ $r_c^*, r_t^*, K_1^* \leftarrow \{0, 1\}^k$,” and we do not use α^* anymore.

Game 6: Same as Game 5, except that the key/commitment pair (ck, tag^*) and the key/ciphertext pair (pk, c^*) and a punctured secret key $\widehat{sk}_{\text{tag}^*}$ are sampled obliviously, and correspondingly the randomness \widehat{r}_c and \widehat{r}_t used for oblivious sampling are used in $r_{\mathcal{A}'}$.

More precisely, the steps “ $r_g, r_c \leftarrow \{0, 1\}^k$; $ck \leftarrow \text{CKG}(1^k; r_g)$; $\text{tag}^* \leftarrow \text{Com}(ck, (c_{\text{in}0}^* \| c_{\text{in}1}^*); r_c)$; $\widehat{r}_c \leftarrow \text{rSamp}_{\mathcal{C}}(r_g, r_c, (c_{\text{in}0}^* \| c_{\text{in}1}^*))$ ” in Game 5 are replaced with the steps “ $\widehat{r}_c \leftarrow \{0, 1\}^k$; $(ck, \text{tag}^*) \leftarrow \text{oSamp}_{\mathcal{C}}(1^k; \widehat{r}_c)$ ”.

Furthermore, the steps “ $r'_g, r'_t \leftarrow \{0, 1\}^k$; $(pk, sk) \leftarrow \text{TKG}(1^k; r'_g)$; $c^* \leftarrow \text{TEnc}(pk, \text{tag}^*, (c_{\text{in}0}^* \| c_{\text{in}1}^*); r'_t)$; $\widehat{r}_t \leftarrow \text{rSamp}_{\mathcal{T}}(r'_g, r'_t, \text{tag}^*, (c_{\text{in}0}^* \| c_{\text{in}1}^*))$ ” in Game 5 are replaced with the steps “ $\widehat{r}_t \leftarrow \{0, 1\}^k$; $(pk, \widehat{sk}_{\text{tag}^*}, c^*) \leftarrow \text{oSamp}_{\mathcal{T}}(\text{tag}^*; \widehat{r}_t)$ ”.

The above completes the description of the games.

For $i \in [5]$, let Succ_i denote the event that \mathcal{A} succeeds in guessing the challenge bit (i.e. $b' = b$ occurs) in Game i . Furthermore, for $i \in \{3, \dots, 6\}$, we define the following *bad* events in Game i :

Bad_i : \mathcal{A} submits a decapsulation query $C = (\text{tag}, c)$ satisfying the following conditions simultaneously: (1) $\text{tag} \neq \text{tag}^*$, (2) $\widehat{\text{Decap}}(\widehat{sk}_{\text{tag}^*}, \text{tag}, c) = (c_{\text{in}0} \| c_{\text{in}1}) \neq \perp$, and (3) $\text{Decap}_{\text{in}}(sk_{\text{in}0}, c_{\text{in}0}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, (1, c_{\text{in}0}))$ or $\text{Decap}_{\text{in}}(sk_{\text{in}1}, c_{\text{in}1}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, (2, c_{\text{in}1}))$.

$\text{Bad}_i^{(\sigma)}$: (where $\sigma \in \{0, 1\}$) \mathcal{A} submits a decapsulation query $C = (\text{tag}, c)$ that satisfies the same conditions as Bad_i , except that the condition (3) is replaced with the condition: $\text{Decap}_{\text{in}}(sk_{\text{in}\sigma}, c_{\text{in}\sigma}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, (\sigma + 1, c_{\text{in}\sigma}))$.

Bad_i^* : \mathcal{A} submits a decapsulation query $C = (\text{tag}, c)$ that satisfies the same conditions as Bad_i , except that the condition (3) is replaced with the condition: $\text{Decap}_{\text{in}}(sk_{\text{in}\gamma}, c_{\text{in}\gamma}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, (\gamma + 1, c_{\text{in}\gamma}))$ (where γ is the random bit chosen just before executing \mathcal{A}).

Note that for all $i \in \{3, \dots, 6\}$, the events $\text{Bad}_i^{(0)}$, $\text{Bad}_i^{(1)}$, and Bad_i^* all imply the event Bad_i , and thus we have $\Pr[\text{Bad}_i^{(0)}], \Pr[\text{Bad}_i^{(1)}], \Pr[\text{Bad}_i^*] \leq \Pr[\text{Bad}_i]$.

By the definitions of the games and events, we have

$$\begin{aligned} \text{Adv}_{\Gamma, \mathcal{A}}^{\text{CCA}}(k) &= 2 \cdot \left| \Pr[\text{Succ}_1] - \frac{1}{2} \right| \\ &\leq 2 \cdot \left(\sum_{i \in [4]} \left| \Pr[\text{Succ}_i] - \Pr[\text{Succ}_{i+1}] \right| + \left| \Pr[\text{Succ}_5] - \frac{1}{2} \right| \right). \quad (1) \end{aligned}$$

In the following, we will upperbound each term that appears in the right hand side of the above inequality.

Claim 1 *There exists a PPTA \mathcal{B}_b such that $\text{Adv}_{\mathcal{C}, \mathcal{B}_b}^{\text{TBind}}(k) \geq |\Pr[\text{Succ}_1] - \Pr[\text{Succ}_2]|$.*

Proof of Claim 1. For $i \in \{1, 2\}$, let NoBind_i be the event that in Game i , \mathcal{A} submits at least one decapsulation query $C = (\text{tag}, c)$ satisfying $\text{tag} = \text{tag}^*$ and $\text{Decap}(SK, C) \neq \perp$. Recall that \mathcal{A} 's query C must satisfy $C \neq C^* = (\text{tag}^*, c^*)$, and thus $\text{tag} = \text{tag}^*$ implies $c \neq c^*$. The difference between Game 1 and Game 2 is how \mathcal{A} 's decapsulation query $C = (\text{tag}, c)$ satisfying $\text{tag} = \text{tag}^*$ is answered. Hence, these games proceed identically unless NoBind_1 or NoBind_2 occurs in the corresponding games, and thus we have

$$\left| \Pr[\text{Succ}_1] - \Pr[\text{Succ}_2] \right| \leq \Pr[\text{NoBind}_1] = \Pr[\text{NoBind}_2]. \quad (2)$$

Thus, it is sufficient to upperbound $\Pr[\text{NoBind}_2]$.

Observe that for a decapsulation query $C = (\text{tag}^*, c)$ satisfying the condition of NoBind_2 , it is guaranteed that $\text{TDec}(sk, \text{tag}, c) = (c_{\text{in}0} \| c_{\text{in}1}) \neq (c_{\text{in}0}^* \| c_{\text{in}1}^*)$. Indeed, if $\text{TDec}(sk, \text{tag}, c) = (c_{\text{in}0}^* \| c_{\text{in}1}^*)$ and $\text{Decap}(SK, C) \neq \perp$, then by the validity check of c in Decap , we have $c^* = c$, which is because c must satisfy $\text{TEnc}(pk, \text{tag}^*, (c_{\text{in}0}^* \| c_{\text{in}1}^*); r_t^*) = c$ where r_t^* is the $(k+1)$ -to- $2k$ -th bits of $\alpha^* = (\alpha_0^* \oplus \alpha_1^*) = (\text{Decap}_{\text{in}}(sk_{\text{in}0}, c_{\text{in}0}^*) \oplus \text{Decap}_{\text{in}}(sk_{\text{in}1}, c_{\text{in}1}^*))$. However, $\text{TEnc}(pk, \text{tag}^*, (c_{\text{in}0}^* \| c_{\text{in}1}^*); r_t^*) = c^*$ also holds due to how c^* is generated, and thus contradicting the condition $c \neq c^*$ implied by NoBind_2 .

We use the above fact to show how to construct a PPTA adversary \mathcal{B}_b that attacks the target-binding property of the commitment scheme \mathcal{C} with advantage $\text{Adv}_{\mathcal{C}, \mathcal{B}_b}^{\text{TBind}}(k) = \Pr[\text{NoBind}_2]$. The description of $\mathcal{B}_b = (\mathcal{B}_{b1}, \mathcal{B}_{b2})$ is as follows:

$\mathcal{B}_{b1}(1^k)$: \mathcal{B}_{b1} first runs $(pk_{\text{in}0}, sk_{\text{in}0}) \leftarrow \text{KKG}_{\text{in}}(1^k)$, $(pk_{\text{in}1}, sk_{\text{in}1}) \leftarrow \text{KKG}_{\text{in}}(1^k)$, $(c_{\text{in}0}^*, \alpha_0^*) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}0})$, and $(c_{\text{in}1}^*, \alpha_1^*) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}1})$. \mathcal{B}_{b1} then sets $\alpha^* \leftarrow (\alpha_0^* \oplus \alpha_1^*)$, and parses α^* as $(r_c^*, r_t^*, \alpha^*) \in (\{0, 1\}^k)^3$. Finally, \mathcal{B}_{b1} sets $M \leftarrow (c_{\text{in}0}^* \| c_{\text{in}1}^*)$, $R \leftarrow r_c^*$, and $\text{st}_{\mathcal{B}} \leftarrow (\mathcal{B}_{b1}$'s entire view), and terminates with output $(M, R, \text{st}_{\mathcal{B}})$.

$\mathcal{B}_{b2}(\text{st}_{\mathcal{B}}, ck)$: \mathcal{B}_{b2} first runs $(pk, sk) \leftarrow \text{TKG}(1^k)$, and then sets $PK \leftarrow (pk_{\text{in}0}, pk_{\text{in}1}, pk, ck)$ and $SK \leftarrow (sk_{\text{in}0}, sk_{\text{in}1}, sk, PK)$. \mathcal{B}_{b2} next runs $\text{tag}^* \leftarrow \text{Com}(ck, (c_{\text{in}0}^* \| c_{\text{in}1}^*); r_c^*)$ and $c^* \leftarrow \text{TEnc}(pk, \text{tag}^*, (c_{\text{in}0}^* \| c_{\text{in}1}^*); r_t^*)$, sets $C^* \leftarrow (\text{tag}^*, c^*)$, and also chooses $K_0^* \in \{0, 1\}^k$ and $b \in \{0, 1\}$ uniformly at random. Then, \mathcal{B}_{b2} runs \mathcal{A} , where the decapsulation queries from \mathcal{A} are answered as Game 2 does, which is possible because \mathcal{B}_{b2} possesses SK .

When \mathcal{A} terminates, \mathcal{B}_{b2} checks if \mathcal{A} has made a decapsulation query $C = (\text{tag}, c)$ satisfying the conditions of NoBind_2 , namely, $\text{tag} = \text{tag}^*$, $c \neq c^*$, $\text{TDec}(sk, \text{tag}, c) = (c_{\text{in}0} \| c_{\text{in}1}) \notin \{(c_{\text{in}0}^* \| c_{\text{in}1}^*), \perp\}$, $\text{Decap}_{\text{in}}(sk_{\text{in}0}, c_{\text{in}0}) = \alpha_0 \neq \perp$, $\text{Decap}_{\text{in}}(sk_{\text{in}1}, c_{\text{in}1}) = \alpha_1 \neq \perp$, $(\alpha_0 \oplus \alpha_1) = (r_c \| r_t \| K) \in \{0, 1\}^{3k}$, and $\text{Com}(ck, (c_{\text{in}0} \| c_{\text{in}1}); r_c) = \text{tag}^*$, and $\text{TEnc}(pk, \text{tag}, (c_{\text{in}0} \| c_{\text{in}1}); r_t) = c$. (Actually, the last condition is redundant for \mathcal{B}_{b2} 's purpose.) If such a query is found, then \mathcal{B}_{b2} terminates with output $M' = (c_{\text{in}0} \| c_{\text{in}1})$ and $R' = r_c$. Otherwise, \mathcal{B}_{b2} gives up and aborts.

The above completes the description of \mathcal{B}_b . It is easy to see that \mathcal{B}_b does a perfect simulation of Game 2 for \mathcal{A} , and whenever \mathcal{A} makes a query that causes

the event NoBind_2 , \mathcal{B}_{b2} can find such a query by using SK and output a pair $(M', R') = ((c_{\text{in}0} \| c_{\text{in}1}), r_c)$ satisfying $\text{Com}(ck, M; R) = \text{Com}(ck, M'; R') = \text{tag}^*$ and $M \neq M'$, violating the target-binding property of the commitment scheme \mathcal{C} . Therefore, we have $\text{Adv}_{\mathcal{C}, \mathcal{B}_b}^{\text{TBind}}(k) = \Pr[\text{NoBind}_2]$. Then, by Equation (2), we have $\text{Adv}_{\mathcal{C}, \mathcal{B}_b}^{\text{TBind}}(k) \geq |\Pr[\text{Succ}_1] - \Pr[\text{Succ}_2]|$, as required. \square (**Claim 1**)

Claim 2 $\Pr[\text{Succ}_2] = \Pr[\text{Succ}_3]$.

Proof of Claim 2. It is sufficient to show that the behavior of the oracle given to \mathcal{A} in Game 2 and that in Game 3 are identical. Let $C = (\text{tag}, c)$ be a decapsulation query that \mathcal{A} makes. If $\text{tag} = \text{tag}^*$, then the query is answered with \perp in Game 2 by definition, while the oracle $\text{AltDecap}(\widehat{SK}_{\text{tag}^*}, C)$ that is given access to \mathcal{A} in Game 3 also returns \perp by definition. Otherwise (i.e. $\text{tag} \neq \text{tag}^*$), by Lemma 3, the result of $\text{Decap}(SK, C)$ and that of $\text{AltDecap}(\widehat{SK}_{\text{tag}^*}, C)$ always agree. This completes the proof. \square (**Claim 2**)

Claim 3 *There exist PPTAs \mathcal{B}_g and \mathcal{B}_a such that*

$$|\Pr[\text{Succ}_3] - \Pr[\text{Succ}_4]| \leq 2 \cdot \left(\text{Adv}_{\Gamma_{\text{in}}, \mathcal{B}_g}^{\text{CPA}}(k) + \text{Adv}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_a}^{\text{TS}}(k) + \text{Adv}_{\Gamma_{\text{in}}, \mathcal{A}', \varepsilon, 2}^{\text{SPA1}}(k) \right).$$

We postpone the proof of this claim to the end of the proof of Theorem 1.

Claim 4 *There exists a PPTA \mathcal{B}'_g such that $\text{Adv}_{\Gamma_{\text{in}}, \mathcal{B}'_g}^{\text{CPA}}(k) = |\Pr[\text{Succ}_4] - \Pr[\text{Succ}_5]|$.*

Proof of Claim 2. Using \mathcal{A} and \mathcal{E} as building blocks, we show how to construct a PPTA CPA adversary \mathcal{B}'_g with the claimed advantage. The description of \mathcal{B}'_g is as follows:

$\mathcal{B}'_g(pk', c^*, \alpha_\beta^*)$: (where $\beta \in \{0, 1\}$ is \mathcal{B}'_g 's challenge bit in its CPA experiment) \mathcal{B}'_g sets $pk_{\text{in}0} \leftarrow pk'$, $c_{\text{in}0}^* \leftarrow c^*$, and $\alpha_0^* \leftarrow \alpha_\beta^*$. Next, \mathcal{B}'_g generates $(pk_{\text{in}1}, sk_{\text{in}1}) \leftarrow \text{KKG}_{\text{in}}(1^k)$ and $(c_{\text{in}1}^*, \alpha_1^*) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}1})$, sets $\alpha^* \leftarrow (\alpha_0^* \oplus \alpha_1^*)$, and parses α^* as $(r_c^*, r_t^*, K_1^*) \in (\{0, 1\}^k)^3$. Then, \mathcal{B}'_g picks $r_g, r'_g \leftarrow \{0, 1\}^*$ uniformly at random, and runs $ck \leftarrow \text{CKG}(1^k; r_g)$, $\text{tag}^* \leftarrow \text{Com}(ck, (c_{\text{in}0}^* \| c_{\text{in}1}^*); r_c^*)$, $\widehat{r}_c \leftarrow \text{rSamp}_{\mathcal{C}}(r_g, r_c^*, (c_{\text{in}0}^* \| c_{\text{in}1}^*))$, $(pk, sk) \leftarrow \text{TKG}(1^k; r'_g)$, $\widehat{sk}_{\text{tag}^*} \leftarrow \text{Punc}(sk, \text{tag}^*)$, $c^* \leftarrow \text{TEnc}(pk, \text{tag}^*, (c_{\text{in}0}^* \| c_{\text{in}1}^*); r_t^*)$, and $\widehat{r}_t \leftarrow \text{rSamp}_{\mathcal{T}}(r'_g, r_t^*, \text{tag}^*, (c_{\text{in}0}^* \| c_{\text{in}1}^*))$. Then \mathcal{B}'_g picks $r_{\mathcal{A}} \in \{0, 1\}^*$, $K_0^* \in \{0, 1\}^k$, and $b \in \{0, 1\}$ all uniformly at random, and sets $PK \leftarrow (pk_{\text{in}0}, pk_{\text{in}1}, pk, ck)$, $C^* \leftarrow (\text{tag}^*, c^*)$, $r_{\mathcal{A}'} \leftarrow (r_{\mathcal{A}}, \widehat{r}_c, \widehat{r}_t, K_b^*)$, and $\text{st}_{\mathcal{E}} \leftarrow (pk_{\text{in}0}, pk_{\text{in}1}, r_{\mathcal{A}'})$. Finally, \mathcal{B}'_g runs $\mathcal{A}(PK, C^*, K_b^*; r_{\mathcal{A}})$.

\mathcal{B}'_g answers \mathcal{A} 's decapsulation queries as $\text{AltDecap}'_{\mathcal{E}}$ does, where the initial state of $\text{AltDecap}'_{\mathcal{E}}$ is tag^* , $\widehat{sk}_{\text{tag}^*}$, and $\text{st}_{\mathcal{E}}$. (Note that $\text{st}_{\mathcal{E}}$ is used by \mathcal{E} , and may be updated upon each call of $\text{AltDecap}'_{\mathcal{E}}$.)

When \mathcal{A} terminates with output b' , \mathcal{B}'_g sets $\beta' \leftarrow (b' \stackrel{?}{=} b)$, and terminates with output β' .

The above completes the description of \mathcal{B}'_g . \mathcal{B}'_g 's CPA advantage can be calculated as follows:

$$\begin{aligned} \text{Adv}_{\Gamma_{\text{in}}, \mathcal{B}'_g}^{\text{CPA}}(k) &= 2 \cdot \left| \Pr[\beta' = \beta] - \frac{1}{2} \right| = \left| \Pr[\beta' = 1 | \beta = 1] - \Pr[\beta' = 1 | \beta = 0] \right| \\ &= \left| \Pr[b' = b | \beta = 1] - \Pr[b' = b | \beta = 0] \right|. \end{aligned}$$

Consider the case when $\beta = 1$. It is easy to see that in this case, \mathcal{B}'_g simulates Game 4 perfectly for \mathcal{A} . Specifically, the real session-key $\alpha'_{\beta^*} = \alpha_1^*$ (corresponding to $c_{\text{in}0}^* = c^*$) is used as α_0^* , and thus $\alpha^* = (\alpha_0^* \oplus \alpha_1^*) = (r_c^* \| r_t^* \| K_1^*)$ is generated exactly as that in Game 4. All other values are distributed identically to those in Game 4. Furthermore, \mathcal{B}'_g uses $\text{AltDecap}'_{\mathcal{E}}$ for answering \mathcal{A} 's decapsulation queries, where the initial state of $\text{AltDecap}'_{\mathcal{E}}$ (and thus the initial state of \mathcal{E}) is appropriately generated as those in Game 4. Under this situation, the probability that \mathcal{A} succeeds in guessing b (i.e. $b' = b$ occurs) is exactly the same as the probability that \mathcal{A} does so in Game 4, i.e. $\Pr[b' = b | \beta = 1] = \Pr[\text{Succ}_4]$.

On the other hand, when $\beta = 0$, then \mathcal{B}'_g simulates Game 5 perfectly for \mathcal{A} . Specifically, in this case, a uniformly random value $\alpha'_{\beta^*} = \alpha_0^*$ is used as α_0^* . Therefore, $\alpha^* = (\alpha_0^* \oplus \alpha_1^*)$ is also a uniformly random $3k$ -bit string, and thus each of r_c^* , r_t^* , and K_1^* is a uniformly random k -bit string, which is exactly how these values are chosen in Game 5. Since this is the only change from the case of $\beta = 1$, with a similar argument to the above, we have $\Pr[b' = b | \beta = 0] = \Pr[\text{Succ}_5]$.

In summary, we have $\text{Adv}_{\Gamma_{\text{in}}, \mathcal{B}'_g}^{\text{CPA}}(k) = |\Pr[\text{Succ}_4] - \Pr[\text{Succ}_5]|$, as required.

□ (Claim 4)

Claim 5 $\Pr[\text{Succ}_5] = 1/2$.

Proof of Claim 5. This is obvious because in Game 5, the real session-key K_1^* is made independent of the challenge ciphertext C^* . Since both K_1^* and K_0^* are now uniformly random, the view of \mathcal{A} does not contain any information on b . This means that the probability that \mathcal{A} succeeds in guessing the challenge bit is exactly $1/2$. □ (Claim 5)

Claims 1 to 5 and Equation (1) guarantee that there exist PPTAs \mathcal{B}_b , \mathcal{B}_g , \mathcal{B}_a , and \mathcal{B}'_g such that

$$\begin{aligned} \text{Adv}_{\Gamma, \mathcal{A}}^{\text{CCA}}(k) &\leq 2 \cdot \text{Adv}_{\mathcal{C}, \mathcal{B}_b}^{\text{TBind}}(k) + 4 \cdot \text{Adv}_{\Gamma_{\text{in}}, \mathcal{B}_g}^{\text{CPA}}(k) + 4 \cdot \text{Adv}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_a}^{\text{TS}}(k) \\ &\quad + 4 \cdot \text{Adv}_{\Gamma_{\text{in}}, \mathcal{A}', \mathcal{E}, 2}^{\text{sPA1}}(k) + 2 \cdot \text{Adv}_{\Gamma_{\text{in}}, \mathcal{B}'_g}^{\text{CPA}}(k), \end{aligned}$$

which, due to our assumptions on the building blocks and Lemma 2, implies that $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{CCA}}(k)$ is negligible. Recall that the choice of the PPTA CCA adversary \mathcal{A} was arbitrarily, and thus for any PPTA CCA adversary \mathcal{A} we can show a negligible upperbound for $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{CCA}}(k)$ as above.

In order to finish the proof of Theorem 1, it remains to prove Claim 3.

Proof of Claim 3. Note that the difference between Game 3 and Game 4 is how a query $C = (\text{tag}, c)$ satisfying the conditions of Bad_3 (or Bad_4) is answered,

and Game 3 and Game 4 proceed identically unless Bad_3 or Bad_4 occurs in the corresponding games. This means that we have

$$\left| \Pr[\text{Succ}_3] - \Pr[\text{Succ}_4] \right| \leq \Pr[\text{Bad}_3] = \Pr[\text{Bad}_4]. \quad (3)$$

We claim the following:

Subclaim 1 $\Pr[\text{Bad}_4] \leq 2 \cdot \Pr[\text{Bad}_4^*]$.

Proof of Subclaim 1. The argument here is essentially the same as the one used in the proof of Claim 4.13 in [17].

Note that the event Bad_4 , $\text{Bad}_4^{(0)}$, $\text{Bad}_4^{(1)}$, and Bad_4^* are triggered once \mathcal{A} makes a query $C = (\text{tag}, c)$ satisfying the conditions that cause these events. Moreover, by definition, if any of the latter three events occurs, then Bad_4 occurs. Furthermore, the bit γ is information-theoretically hidden from \mathcal{A} 's view in Game 4. This means that the probability of Bad_4^* occurring is identical to the probability of the event (in Game 4) that is triggered when (1) \mathcal{A} first makes a query satisfying the conditions of Bad_4 , (2) γ is picked ‘‘on-the-fly’’ at this point, and then (3) $\text{Decap}_{\text{in}}(sk_{\text{in}\gamma}, c_{\text{in}\gamma}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, (\gamma + 1, c_{\text{in}\gamma}))$ holds. The probability of this event occurring is $\Pr_{\gamma \leftarrow \{0,1\}}[\text{Bad}_4 \wedge \text{Bad}_4^{(\gamma)}] = \Pr_{\gamma \leftarrow \{0,1\}}[\text{Bad}_4^{(\gamma)}]$ (where the probability is also over Game 4 except the choice of γ). This can be further estimated as follows:

$$\begin{aligned} \Pr_{\gamma \leftarrow \{0,1\}}[\text{Bad}_4^{(\gamma)}] &= \frac{1}{2} \left(\Pr[\text{Bad}_4^{(0)}] + \Pr[\text{Bad}_4^{(1)}] \right) \\ &\geq \frac{1}{2} \Pr[\text{Bad}_4^{(0)} \vee \text{Bad}_4^{(1)}] = \frac{1}{2} \Pr[\text{Bad}_4], \end{aligned}$$

where we used $\Pr[\text{Bad}_4^{(0)} \vee \text{Bad}_4^{(1)}] = \Pr[\text{Bad}_4]$, which is by definition.

In summary, we have $\Pr[\text{Bad}_4^*] \geq \frac{1}{2} \Pr[\text{Bad}_4]$, as required. \square (**Subclaim 1**)

Using Subclaim 1, we can further estimate $\Pr[\text{Bad}_4]$ as follows:

$$\begin{aligned} \Pr[\text{Bad}_4] &\leq 2 \cdot \Pr[\text{Bad}_4^*] \\ &\leq 2 \cdot \left(\left| \Pr[\text{Bad}_4^*] - \Pr[\text{Bad}_5^*] \right| + \Pr[\text{Bad}_5^*] \right) \\ &\leq 2 \cdot \left(\left| \Pr[\text{Bad}_4^*] - \Pr[\text{Bad}_5^*] \right| + \Pr[\text{Bad}_5] \right) \\ &\leq 2 \cdot \left(\left| \Pr[\text{Bad}_4^*] - \Pr[\text{Bad}_5^*] \right| + \left| \Pr[\text{Bad}_5] - \Pr[\text{Bad}_6] \right| + \Pr[\text{Bad}_6] \right), \end{aligned} \quad (4)$$

where we used $\Pr[\text{Bad}_5^*] \leq \Pr[\text{Bad}_5]$ in the third inequality, which is again by definition. It remains to upperbound the right hand side of the above inequality.

Subclaim 2 *There exists a PPTA $\mathcal{B}_{\mathbf{g}}$ such that $\text{Adv}_{I_{\text{in}}, \mathcal{B}_{\mathbf{g}}}^{\text{CPA}}(k) = |\Pr[\text{Bad}_4^*] - \Pr[\text{Bad}_5^*]|$.*

Proof of Subclaim 2. Using \mathcal{A} and \mathcal{E} as building blocks, we show how to construct a PPTA CPA adversary \mathcal{B}_g with the claimed advantage. The description of \mathcal{B}_g is as follows:

$\mathcal{B}_g(pk', c^*, \alpha_\beta^*)$: (where $\beta \in \{0, 1\}$ is \mathcal{B}_g 's challenge bit in its CPA experiment) \mathcal{B}_g picks $\gamma \in \{0, 1\}$ uniformly at random, then sets $pk_{\text{in}(1-\gamma)} \leftarrow pk'$, $c_{\text{in}(1-\gamma)}^* \leftarrow c^*$, and $\alpha_{1-\gamma}^* \leftarrow \alpha_\beta^*$. Next, \mathcal{B}_g generates $(pk_{\text{in}\gamma}, sk_{\text{in}\gamma}) \leftarrow \text{KKG}_{\text{in}}(1^k)$ and $(c_{\text{in}\gamma}^*, \alpha_\gamma^*) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}\gamma})$, sets $\alpha^* \leftarrow (\alpha_0^* \oplus \alpha_1^*)$, and parses α^* as $(r_c^*, r_t^*, K_1^*) \in (\{0, 1\}^k)^3$. Then, \mathcal{B}_g prepares $K_1^*, K_0^* \in \{0, 1\}^k$, $b \in \{0, 1\}$, $PK = (pk_{\text{in}0}, pk_{\text{in}1}, pk, c)$, $C^* = (\text{tag}^*, c^*)$, $\widehat{sk}_{\text{tag}^*}$, and $\text{st}_{\mathcal{E}} = (pk_{\text{in}0}, pk_{\text{in}1}, r_{\mathcal{A}'} = (r_{\mathcal{A}}, \widehat{r}_c, \widehat{r}_t, K_b^*))$, exactly as \mathcal{B}'_g in the proof of Claim 4 does. Finally, \mathcal{B}_g runs $\mathcal{A}(PK, C^*, K_b^*; r_{\mathcal{A}'})$ until it terminates, where \mathcal{B}_g answers \mathcal{A} 's queries in exactly the same way as \mathcal{B}'_g does.

When \mathcal{A} terminates, \mathcal{B}_g checks whether \mathcal{A} has submitted a decapsulation query $C = (\text{tag}, c)$ that satisfies the conditions of Bad_4^* (i.e. (1) $\text{tag} \neq \text{tag}^*$, (2) $\widehat{\text{TDec}}(\widehat{sk}_{\text{tag}^*}, \text{tag}, c) = (c_{\text{in}0} \| c_{\text{in}1}) \neq \perp$, and (3) $\text{Decap}_{\text{in}}(sk_{\text{in}\gamma}, c_{\text{in}\gamma}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, c_{\text{in}\gamma})$ hold), which can be checked by using $sk_{\text{in}\gamma}$. If such a query is found, the \mathcal{B}_g sets $\beta' \leftarrow 1$, otherwise sets $\beta' \leftarrow 0$, and terminates with output β' .

The above completes the description of \mathcal{B}_g . Let $\text{Bad}_{\mathcal{B}}^*$ be the event that \mathcal{A} submits a decapsulation query that satisfies the conditions (1), (2), and (3) of Bad_4^* , in the experiment simulated by \mathcal{B}_g . Note that \mathcal{B}_g outputs $\beta' = 1$ only when $\text{Bad}_{\mathcal{B}}^*$ occurs. Therefore, \mathcal{B}_g 's CPA advantage can be calculated as follows:

$$\begin{aligned} \text{Adv}_{I_{\text{in}}, \mathcal{B}_g}^{\text{CPA}}(k) &= 2 \cdot \left| \Pr[\beta' = \beta] - \frac{1}{2} \right| = \left| \Pr[\beta' = 1 | \beta = 1] - \Pr[\beta' = 1 | \beta = 0] \right| \\ &= \left| \Pr[\text{Bad}_{\mathcal{B}}^* | \beta = 1] - \Pr[\text{Bad}_{\mathcal{B}}^* | \beta = 0] \right|. \end{aligned}$$

With essentially the same arguments as in the proof of Claim 4, we can see that \mathcal{B}_g does a perfect simulation of Game 4 for \mathcal{A} if $\beta = 1$, and does a perfect simulation of Game 5 for \mathcal{A} if $\beta = 0$. In particular, the only difference from the proof of Claim 4 is in which of the positions $(pk_{\text{in}0}, c_{\text{in}0}^*, \alpha_0^*)$ or $(pk_{\text{in}1}, c_{\text{in}1}^*, \alpha_1^*)$ \mathcal{B}_g embeds \mathcal{B}_g 's instance of the CPA experiment. In the proof of Claim 4, the reduction algorithm \mathcal{B}'_g embeds its challenge into $(pk_{\text{in}0}, c_{\text{in}0}^*, \alpha_0^*)$, while in the current proof, the reduction algorithm \mathcal{B}_g embeds its challenge into $(pk_{\text{in}(1-\gamma)}, c_{\text{in}(1-\gamma)}^*, \alpha_{1-\gamma}^*)$ for a random $\gamma \in \{0, 1\}$. It is easy to see that even after this change, if $\beta = 1$, then the view of \mathcal{A} is identical to that in Game 4, and if $\beta = 0$, then the view of \mathcal{A} is identical to that in Game 5.

Under the situation, the probability that $\text{Bad}_{\mathcal{B}}^*$ occurs in the experiment simulated by \mathcal{B}_g in case $\beta = 1$ (resp. $\beta = 0$) is identical to the probability that Bad_4^* (resp. Bad_5^*) occurs in Game 4 (resp. Game 5), namely, we have $\Pr[\text{Bad}_{\mathcal{B}}^* | \beta = 1] = \Pr[\text{Bad}_4^*]$ and $\Pr[\text{Bad}_{\mathcal{B}}^* | \beta = 0] = \Pr[\text{Bad}_5^*]$.

In summary, we have $\text{Adv}_{I_{\text{in}}, \mathcal{B}_g}^{\text{CPA}}(k) = |\Pr[\text{Bad}_4^*] - \Pr[\text{Bad}_5^*]|$, as required.

□ (**Subclaim 2**)

Subclaim 3 *There exists a PPTA \mathcal{B}_d such that $\text{Adv}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS}}(k) = |\text{Pr}[\text{Bad}_5] - \text{Pr}[\text{Bad}_6]|$.*

Proof of Subclaim 3. Using \mathcal{A} and \mathcal{E} as building blocks, we show how to construct a PPTA \mathcal{B} that has the claimed advantage in distinguishing the distributions considered in Lemma 2. The description of $\mathcal{B}_d = (\mathcal{B}_{d1}, \mathcal{B}_{d2})$ as follows:

$\mathcal{B}_{d1}(1^k)$: \mathcal{B}_{d1} runs $(pk_{\text{in}0}, sk_{\text{in}0}) \leftarrow \text{KKG}_{\text{in}}(1^k)$, $(pk_{\text{in}1}, sk_{\text{in}1}) \leftarrow \text{KKG}_{\text{in}}(1^k)$, $(c_{\text{in}0}^*, \alpha_0^*) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}0})$, $(c_{\text{in}1}^*, \alpha_1^*) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}1})$. Then \mathcal{B}_{d1} sets $M \leftarrow (c_{\text{in}0}^* \| c_{\text{in}1}^*)$ and $\text{st}_{\mathcal{B}} \leftarrow (\mathcal{B}_{d1}$'s entire view), and terminates with output $(M, \text{st}_{\mathcal{B}})$.

$\mathcal{B}_{d2}(\text{st}_{\mathcal{B}}, ck, \text{tag}^*, pk, c^*, \widehat{sk}_{\text{tag}^*}, \widehat{r}_c, \widehat{r}_t)$: \mathcal{B}_{d2} sets $PK \leftarrow (pk_{\text{in}0}, pk_{\text{in}1}, pk, ck)$ and $C^* \leftarrow (\text{tag}^*, c^*)$, picks $K^* \in \{0, 1\}^*$ and $r_{\mathcal{A}} \in \{0, 1\}^*$ uniformly at random, and then sets $r_{\mathcal{A}'} \leftarrow (r_{\mathcal{A}}, \widehat{r}_c, \widehat{r}_t, K^*)$ and $\text{st}_{\mathcal{E}} \leftarrow (pk_{\text{in}0}, pk_{\text{in}1}, r_{\mathcal{A}'})$. (Recall that K_0^* and K_1^* in Games 5 and 6 are distributed identically, and thus it is sufficient to choose just a single value K^* and pretend as if K^* is K_b^* .) Then \mathcal{B}_{d2} runs $\mathcal{A}(PK, C^*, K^*; r_{\mathcal{A}})$.

\mathcal{B}_{d2} answers \mathcal{A} 's queries as Game 5 does, which is possible because \mathcal{B}_{d2} possesses $\widehat{sk}_{\text{tag}^*}$ and $\text{st}_{\mathcal{E}}$, and thus \mathcal{B}_{d2} can run $\text{AltDecap}'_{\mathcal{E}}$ (which internally runs the extractor $\mathcal{E}(\text{st}_{\mathcal{E}}, \cdot)$).

When \mathcal{A} terminates, \mathcal{B}_{d2} checks whether \mathcal{A} has submitted a query that satisfies the conditions of Bad_5 , which can be checked by using $sk_{\text{in}0}$ and $sk_{\text{in}1}$ that \mathcal{B}_{d2} possesses. If such a query is found, then \mathcal{B}_{d2} outputs 1, otherwise outputs 0, and terminates.

The above completes the description of \mathcal{B}_d . Let $\text{Bad}_{\mathcal{B}}$ be the event that \mathcal{A} submits a decapsulation query $C = (\text{tag}, c)$ that satisfies the conditions of Bad_5 in the experiment simulated by \mathcal{B}_d (i.e. the query satisfying (1) $\text{tag} \neq \text{tag}^*$, (2) $\widehat{\text{TD}}\widehat{\text{Dec}}(\widehat{sk}_{\text{tag}^*}, \text{tag}, c) = (c_{\text{in}0} \| c_{\text{in}1}) \neq \perp$, and (3) $\text{Decap}_{\text{in}}(sk_{\text{in}0}, c_{\text{in}0}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, c_{\text{in}0})$ or $\text{Decap}_{\text{in}}(sk_{\text{in}1}, c_{\text{in}1}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, c_{\text{in}1})$). Note that \mathcal{B}_d submits 1 only when $\text{Bad}_{\mathcal{B}}$ occurs. Therefore, \mathcal{B}_d 's advantage $\text{Adv}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS}}(k)$ can be calculated as follows:

$$\begin{aligned} \text{Adv}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS}}(k) &= \left| \text{Pr}[\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS-Real}}(k) = 1] - \text{Pr}[\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS-Sim}}(k) = 1] \right| \\ &= \left| \text{Pr}[\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS-Real}} : \text{Bad}_{\mathcal{B}}] - \text{Pr}[\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS-Sim}}(k) : \text{Bad}_{\mathcal{B}}] \right|. \end{aligned}$$

Consider the case when \mathcal{B}_d is run in the “real” experiment $\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS-Real}}(k)$. It is easy to see that in this case, \mathcal{B}_d simulates Game 5 perfectly for \mathcal{A} . Specifically, $ck, pk, \text{tag}^*, c^*$, and $\widehat{sk}_{\text{tag}^*}$ are generated from CKG, TKG, Com, TEnc, and Punc, respectively, in such a way that tag^* is a commitment of $(c_{\text{in}0}^* \| c_{\text{in}1}^*)$ and c^* is an encryption of $(c_{\text{in}0}^* \| c_{\text{in}1}^*)$ under the tag tag^* . Furthermore, \widehat{r}_c and \widehat{r}_t are generated from $\text{rSamp}_{\mathcal{C}}$ and $\text{rSamp}_{\mathcal{T}}$, respectively, which is how they are generated in Game 5. Under the situation, the probability that \mathcal{A} submits a decapsulation query that causes the event $\text{Bad}_{\mathcal{B}}$ is exactly the same as the probability that \mathcal{A} does so in Game 5. That is, we have $\text{Pr}[\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS-Real}}(k) : \text{Bad}_{\mathcal{B}}] = \text{Pr}[\text{Bad}_5]$.

On the other hand, consider the case when \mathcal{B}_d is run in the “simulated” experiment $\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS-Sim}}(k)$. In this case, \mathcal{B}_d simulates Game 6 perfectly for \mathcal{A} .

Specifically, (ck, tag^*) and $(pk, c^*, \widehat{sk}_{\text{tag}^*})$ are generated by $\text{oSamp}_{\mathcal{C}}(1^k; \widehat{r}_c)$ and $\text{oSamp}_{\mathcal{T}}(\text{tag}^*; \widehat{r}_t)$ with uniformly chosen randomness \widehat{r}_c and \widehat{r}_t , respectively, and this is exactly how these values are generated in Game 6. Since this is the only change from the above case, with a similar argument we have $\Pr[\text{Expt}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS-Sim}}(k) : \text{Bad}_{\mathcal{B}}] = \Pr[\text{Bad}_6]$.

In summary, we have $\text{Adv}_{[\mathcal{C}, \mathcal{T}], \mathcal{B}_d}^{\text{TS}}(k) = |\Pr[\text{Bad}_5] - \Pr[\text{Bad}_6]|$, as required. \square (**Subclaim 3**)

Subclaim 4 $\text{Adv}_{\Gamma_{\text{in}}, \mathcal{A}', \mathcal{E}, 2}^{\text{sPA1}}(k) = \Pr[\text{Bad}_6]$.

Proof of Subclaim 4. Note that the view of \mathcal{A} in Game 6 is exactly the same as the view of \mathcal{A} when it is internally run by \mathcal{A}' in the situation where \mathcal{A}' is run in the experiment $\text{Expt}_{\Gamma_{\text{in}}, \mathcal{A}', \mathcal{E}, 2}^{\text{sPA1}}(k)$ with the extractor \mathcal{E} . Therefore, the probability that \mathcal{A} submits a query that causes the event Bad_6 in Game 6, is exactly the same as the probability that \mathcal{A}' submits a query to \mathcal{E} that makes the experiment $\text{Expt}_{\Gamma_{\text{in}}, \mathcal{A}', \mathcal{E}, 2}^{\text{sPA1}}(k)$ outputs 1 (i.e. \mathcal{A}' submits a query of the form $(j+1, c_{\text{in}j})$ such that $\text{Decap}_{\text{in}}(sk_{\text{in}j}, c_{\text{in}j}) \neq \mathcal{E}(\text{st}_{\mathcal{E}}, (j+1, c_{\text{in}j}))$ for some $j \in \{0, 1\}$). \square (**Subclaim 4**)

Equations (3), (4), and Subclaims 2 to 4 imply Claim 3. \square (**Claim 3**)

This concludes the proof of Theorem 1. \square (**Theorem 1**)

4.2 Second Construction

Let $\Gamma_{\text{in}} = (\text{KKG}_{\text{in}}, \text{Encap}_{\text{in}}, \text{Decap}_{\text{in}})$ be a KEM whose ciphertext length is $n = n(k)$ and whose session-key space is $\{0, 1\}^{3k}$ for k -bit security. Let $\mathcal{T} = (\text{TKG}, \text{TEnc}, \text{TDec}, \text{Punc}, \widehat{\text{TDec}})$ be a PTBE scheme and $\mathcal{C} = (\text{CKG}, \text{Com})$ be a commitment scheme. We require the plaintext space of TEnc and the message space of Com to be $\{0, 1\}^n$, and the randomness space of TEnc and that of Com to be $\{0, 1\}^k$ for k -bit security. Then, our second proposed KEM $\overline{\Gamma} = (\text{KKG}, \overline{\text{Encap}}, \overline{\text{Decap}})$ is constructed as in Fig. 4.

The security of $\overline{\Gamma}$ is guaranteed by the following theorem.

Theorem 2. *Assume that the KEM Γ_{in} is 1-CCA secure and sPA1₁ secure, the commitment scheme \mathcal{C} is target-binding and trapdoor simulatable, and the PTBE scheme \mathcal{T} is trapdoor simulatable. Then, the KEM $\overline{\Gamma}$ constructed as in Fig. 4 is CCA secure.*

The proof of this theorem proceeds very similarly to the proof of Theorem 1, and thus we only explain the difference here, and will give the formal proof in the full version.

Recall that in the proof Theorem 1, the “bad” queries (for which the extractor fails to extract correct decapsulation results) are dealt with due to the property of “multiple encryption” of two instances of the KEM Γ_{in} with public keys $(pk_{\text{in}0}, pk_{\text{in}1})$. In particular, the reduction algorithm in the proof of Subclaim 2 that attacks the CPA security of the underlying KEM Γ_{in} , uses one of secret

$\overline{\text{KKG}}(1^k) :$ $(pk_{\text{in}}, sk_{\text{in}}) \leftarrow \text{KKG}_{\text{in}}(1^k)$ $(pk, sk) \leftarrow \text{TKG}(1^k)$ $ck \leftarrow \text{CKG}(1^k)$ $PK \leftarrow (pk_{\text{in}}, pk, ck)$ $SK \leftarrow (sk_{\text{in}}, sk, PK)$ Return (PK, SK) .	$\text{Decap}(SK, C) :$ $(sk_{\text{in}}, sk, PK) \leftarrow SK$ $(pk_{\text{in}}, pk, ck) \leftarrow PK$ $(\text{tag}, c) \leftarrow C$ $c_{\text{in}} \leftarrow \text{TDec}(sk, \text{tag}, c)$ If $c_{\text{in}} = \perp$ then return \perp . $\alpha \leftarrow \text{Decap}_{\text{in}}(sk_{\text{in}}, c_{\text{in}})$ If $\alpha = \perp$ then return \perp . Parse α as $(r_c, r_t, K) \in (\{0, 1\}^k)^3$ If $\text{Com}(ck, c_{\text{in}}; r_c) = \text{tag}$ and $\text{TEnc}(pk, \text{tag}, c_{\text{in}}; r_t) = c$ then return K else return \perp
$\overline{\text{Encap}}(PK) :$ $(pk_{\text{in}}, pk, ck) \leftarrow PK$ $(c_{\text{in}}, \alpha) \leftarrow \text{Encap}_{\text{in}}(pk_{\text{in}})$ Parse α as $(r_c, r_t, K) \in (\{0, 1\}^k)^3$ $\text{tag} \leftarrow \text{Com}(ck, c_{\text{in}}; r_c)$ $c \leftarrow \text{TEnc}(pk, \text{tag}, c_{\text{in}}; r_t)$ $C \leftarrow (\text{tag}, c)$. Return (C, K) .	

Fig. 4. The second proposed construction: the KEM $\overline{\Gamma}$ based on a KEM Γ_{in} , a commitment scheme \mathcal{C} , and a PTBE scheme \mathcal{T} .

keys $sk_{\text{in}\gamma}$ (corresponding to $pk_{\text{in}\gamma}$) to detect whether the bad event occurs, while embedding its CPA instance regarding Γ_{in} into the other position, i.e. into $(pk_{\text{in}(1-\gamma)}, c_{\text{in}(1-\gamma)})$. This strategy works thanks to the argument regarding the probabilities given in the proof of Subclaim 1 (which is in turn based on the proof of [17, Claim 4.13]). However, for this argument to work, it seems to us that we inherently have to rely on the sPA1_2 security of Γ_{in} , in order for the reduction algorithms (especially, the reduction algorithms attacking the CPA of Γ_{in}) to simulate the decapsulation oracle for an adversary \mathcal{A} .

The simple idea employed in our second construction is to change the mechanism of detecting the bad queries by relying on the 1-CCA security of Γ_{in} , so that a reduction algorithm can check (by its access to the decapsulation oracle) whether \mathcal{A} has submitted a bad decapsulation query. This allows us to use Γ_{in} only in the “single” key setting, leading to only requiring it to be sPA1_1 secure. By employing this idea, a security analysis similar to the recent constructions [44, 31, 37, 40] works, and for the other parts of the security proof (other than the analysis regarding dealing with the bad decapsulation queries) are essentially the same as those in the proof of Theorem 1. For more details, see the full version.

On the Merits of the Second Construction. Since we need to use a KEM which simultaneously satisfies 1-CCA and sPA1_1 security for our second construction, a natural question would be whether we can construct such a scheme. We note that we can achieve such a KEM from a CPA secure PKE (or a KEM) which is also sPA1_{2k} secure. Specifically, Dodis and Fiore [21, Appendix C] showed how to construct a 1-CCA secure PKE scheme from the combination of a CPA secure PKE scheme and a one-time secure signature scheme (in which $2k$ independently generated public keys are arranged as in the “DDN-lite” construction, but a message is encoded and encrypted in a k -out-of- k fashion, rather than encrypting

the same message under k public keys). It is straightforward to see that their construction is $\mathsf{sPA1}_1$ secure if the underlying PKE scheme is $\mathsf{sPA1}_{2k}$ secure. We note that we can slightly optimize their construction by using a CPA secure KEM, instead of a PKE scheme, as a building block. We provide the construction and its security proof in the full version.

However, if we implement a 1-CCA and $\mathsf{sPA1}_1$ secure KEM from a CPA and $\mathsf{sPA1}_{2k}$ secure KEM, there is no merit compared to our first construction (that only requires a CPA and $\mathsf{sPA1}_2$ secure KEM), both in terms of the assumptions and the efficiency. So far, we do not know a better way to construct a 1-CCA and $\mathsf{sPA1}_1$ secure scheme than the approach that relies on [21, Appendix C]. We would like to however emphasize that the point of our second construction is that it may in the future be possible to come up with a direct construction of a KEM (or a PKE scheme) satisfying the requirements for the second construction, from assumptions weaker than those required in our first construction or the combination of our second construction and the Dodis-Fiore construction. We believe that such a possibility of the existence of better constructions can be a *raison d'être* of our second construction. In particular, we actually do not need the “full” power of 1-CCA security, but a (seemingly) much weaker security notion such that CPA security holds in the presence of one “plaintext-checking” query [47, 1]. More specifically, a plaintext-checking query (for a KEM it could be called a session-key-checking query, but we stick to the terminology in [47]) is a query of the form (c, K) , and its reply is the one-bit $(\text{Decap}(sk, c) \stackrel{?}{=} K)$. This could be a hint for the next step.

We would also like to note that even if using the result based on [21], we still achieve the property of “separating” the requirement that a single PKE scheme (or a KEM) needs to satisfy “plaintext awareness” and a “simulatability property” simultaneously in [18]. This is another merit of our second construction.

Acknowledgement. The authors would like to thank the members of the study group “Shin-Akarui-Angou-Benkyou-Kai,” and the anonymous reviewers for their helpful comments and suggestions.

References

1. M. Abdalla, F. Benhamouda, and D. Pointcheval. Public-key encryption indistinguishable under plaintext-checkable attacks. *PKC 2015*, LNCS 9020, pp. 332–352, 2015.
2. B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S.P. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. *CRYPTO 2001*, LNCS 2139, pp. 1–18, 2001.
3. M. Bellare, V.T. Hoang, and S. Keelveedhi. Instantiating random oracles via UCEs. *CRYPTO 2013(2)*, LNCS 8043, pp. 398–415, 2013.
4. M. Bellare, D. Hofheinz, and S. Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. *EUROCRYPT 2009*, LNCS 5479, pp. 1–35, 2009.

5. M. Bellare and A. Palacio. Towards plaintext-aware public-key encryption without random oracles. *ASIACRYPT 2004*, LNCS 3329, pp. 48–62, 2004.
6. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. *CCS 1993*, pp. 62–73, 1993.
7. M. Bellare and P. Rogaway. Optimal asymmetric encryption. *EUROCRYPT 1994*, LNCS 950, pp. 92–111, 1995.
8. D. Bleichenbacher. Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. *CRYPTO 1998*, LNCS 1462, pp. 1–12, 1998.
9. R. Canetti. Towards realizing random oracles: Hash functions that hide all partial information. *CRYPTO 1997*, LNCS 1294, pp. 455–469, 1997.
10. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. *FOCS 2001*, pp. 136–145, 2001.
11. R. Canetti, U. Feige, O. Goldreich, and M. Naor. Adaptively secure multi-party computation. *STOC 1996*, pp. 639–648, 1996.
12. R. Canetti, S. Halevi, and J. Katz. Chosen-ciphertext security from identity-based encryption. *EUROCRYPT 2004*, LNCS 3027, pp. 207–222, 2004.
13. Y. Chen and Z. Zhang. Publicly evaluable pseudorandom functions and their applications. *SCN 2014*, LNCS 8642, pp. 115–134, 2014.
14. S.G. Choi, D. Dachman-Soled, T. Malkin, and H. Wee. Improved non-committing encryption with applications to adaptively secure protocols. *ASIACRYPT 2009*, LNCS 5912, pp. 287–302, 2009.
15. R. Cramer, G. Hanaoka, D. Hofheinz, H. Imai, E. Kiltz, R. Pass, A. Shelat, and V. Vaikuntanathan. Bounded CCA2-secure encryption. *ASIACRYPT 2007*, LNCS 4833, pp. 502–518, 2007.
16. R. Cramer and V. Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM J. Computing*, 33(1):167–226, 2003.
17. D. Dachman-Soled. A black-box construction of a CCA2 encryption scheme from a plaintext aware encryption scheme, 2013. Full version of [18]. <http://eprint.iacr.org/2013/680>.
18. D. Dachman-Soled. A black-box construction of a CCA2 encryption scheme from a plaintext aware (SPA1) encryption scheme. *PKC 2014*, LNCS 8383, pp. 37–55, 2014.
19. I. Damgård and J.B. Nielsen. Improved non-committing encryption schemes based on a general complexity assumption. *CRYPTO 2000*, LNCS 1880, pp. 432–450, 2000.
20. A.W. Dent. The Cramer-Shoup encryption is plaintext aware in the standard model. *EUROCRYPT 2006*, LNCS 4004, pp. 289–307, 2006.
21. Y. Dodis and D. Fiore. Interactive encryption and message authentication, 2013. Full version of [22]. <http://eprint.iacr.org/2013/817>.
22. Y. Dodis and D. Fiore. Interactive encryption and message authentication. *SCN 2014*, LNCS 8642, pp. 494–513, 2014.
23. D. Dolev, C. Dwork, and M. Naor. Non-malleable cryptography. *STOC 1991*, pp. 542–552, 1991.
24. E. Fujisaki and T. Okamoto. How to enhance the security of public-key encryption at minimum cost. *PKC 1999*, LNCS 1560, pp. 53–68, 1999.
25. E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *CRYPTO 1999*, LNCS 1666, pp. 537–554, 1999.
26. S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *FOCS 2013*, pp. 40–49, 2013.

27. O. Goldreich and R.D. Rothblum. Enhancements of trapdoor permutations. *J. of Cryptology*, 26(3):484–512, 2013.
28. M. Hajiabadi and B.M. Kapron. Reproducible circularly-secure bit encryption: Applications and realizations. *CRYPTO 2015(1)*, LNCS 9215, pp. 224–243, 2015.
29. B. Hemenway and R. Ostrovsky. On homomorphic encryption and chosen-ciphertext security. *PKC 2012*, LNCS 7293, pp. 52–65, 2012.
30. B. Hemenway and R. Ostrovsky. Building lossy trapdoor functions from lossy encryption. *ASIACRYPT 2013(2)*, LNCS 8270, pp. 241–260, 2013.
31. S. Hohenberger, A. Lewko, and B. Waters. Detecting dangerous queries: A new approach for chosen ciphertext security. *EUROCRYPT 2012*, LNCS 7237, pp. 663–681, 2012.
32. E. Kiltz. Chosen-ciphertext security from tag-based encryption. *TCC 2006*, LNCS 3876, pp. 581–600, 2006.
33. E. Kiltz, P. Mohassel, and A. O’Neill. Adaptive trapdoor functions and chosen-ciphertext security. *EUROCRYPT 2010*, LNCS 6110, pp. 673–692, 2010.
34. H. Lin and S. Tessaro. Amplification of chosen-ciphertext security. *EUROCRYPT 2013*, LNCS 7881, pp. 503–519, 2013.
35. Y. Lindell. A simpler construction of CCA2-secure public-key encryption under general assumptions. *EUROCRYPT 2003*, LNCS 2656, pp. 241–254, 2003.
36. B. Lynn, M. Prabhakaran, and A. Sahai. Positive results and techniques for obfuscation. *EUROCRYPT 2004*, LNCS 3027, pp. 20–39, 2004.
37. T. Matsuda and G. Hanaoka. Achieving chosen ciphertext security from detectable public key encryption efficiently via hybrid encryption. *IWSEC 2013*, LNCS 8231, pp. 226–243, 2013.
38. T. Matsuda and G. Hanaoka. Chosen ciphertext security via point obfuscation. *TCC 2014*, LNCS 8349, pp. 95–120, 2014.
39. T. Matsuda and G. Hanaoka. Chosen ciphertext security via UCE. *PKC 2014*, LNCS 8383, pp. 56–76, 2014.
40. T. Matsuda and G. Hanaoka. An asymptotically optimal method for converting bit encryption to multi-bit encryption. *ASIACRYPT 2015(1)*, LNCS 9452, pp. 415–442, 2015.
41. T. Matsuda and G. Hanaoka. Constructing and understanding chosen ciphertext security via puncturable key encapsulation mechanisms. *TCC 2015(1)*, LNCS 9014, pp. 561–590, 2015.
42. P. Mol and S. Yilek. Chosen-ciphertext security from slightly lossy trapdoor functions. *PKC 2010*, LNCS 6056, pp. 296–311, 2010.
43. S. Myers, M. Sergi, and A. Shelat. Blackbox construction of a more than non-malleable CCA1 encryption scheme from plaintext awareness. *SCN 2012*, LNCS 7485, pp. 149–165, 2012.
44. S. Myers and A. Shelat. Bit encryption is complete. *FOCS 2009*, pp. 607–616, 2009.
45. M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. *STOC 1989*, pp. 33–43, 1989.
46. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. *STOC 1990*, pp. 427–437, 1990.
47. T. Okamoto and D. Pointcheval. REACT: Rapid enhanced-security asymmetric cryptosystem transform. *CT-RSA 2001*, LNCS 2020, pp. 159–174, 2001.
48. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. *STOC 2008*, pp. 187–196, 2008.
49. C. Rackoff and D.R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. *CRYPTO 1991*, LNCS 576, pp. 433–444, 1992.

50. A. Rosen and G. Segev. Chosen-ciphertext security via correlated products. *TCC 2009*, LNCS 5444, pp. 419–436, 2009.
51. A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. *FOCS 1999*, pp. 543–553, 1999.
52. A. Sahai and B. Waters. How to use indistinguishability obfuscation: deniable encryption, and more. *STOC 2014*, pp. 475–484, 2014.
53. A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero-knowledge. *CRYPTO 2001*, LNCS 2139, pp. 566–598, 2001.
54. H. Wee. Efficient chosen-ciphertext security via extractable hash proofs. *CRYPTO 2010*, LNCS 6223, pp. 314–332, 2010.

A Standard Cryptographic Primitives

Public Key Encryption. A public key encryption (PKE) scheme Π consists of the three PPTAs (PKG, Enc, Dec) with the following interface:

$$\begin{array}{lll} \textbf{Key Generation:} & \textbf{Encryption:} & \textbf{Decryption:} \\ \hline (pk, sk) \leftarrow \text{PKG}(1^k) & c \leftarrow \text{Enc}(pk, m) & m \text{ (or } \perp) \leftarrow \text{Dec}(sk, c) \end{array}$$

where Dec is a deterministic algorithm, (pk, sk) is a public/secret key pair, and c is a ciphertext of a plaintext m under pk . We say that a PKE scheme satisfies *correctness* if for all $k \in \mathbb{N}$, all keys (pk, sk) output from $\text{PKG}(1^k)$, and all plaintexts m , it holds that $\text{Dec}(sk, \text{Enc}(pk, m)) = m$.

Since we do not directly use the ordinary security notions for PKE in this paper, we do not introduce them. In Section 2.2, we review the (simplified version of) trapdoor simulatability property [14] of a PKE scheme.

Key Encapsulation Mechanism. A key encapsulation mechanism (KEM) Γ consists of the three PPTAs (KKG, Encap, Decap) with the following interface:

$$\begin{array}{lll} \textbf{Key Generation:} & \textbf{Encapsulation:} & \textbf{Decapsulation:} \\ \hline (pk, sk) \leftarrow \text{KKG}(1^k) & (c, K) \leftarrow \text{Encap}(pk) & K \text{ (or } \perp) \leftarrow \text{Decap}(sk, c) \end{array}$$

where Decap is a deterministic algorithm, (pk, sk) is a public/secret key pair that defines a session-key space \mathcal{K} , and c is a ciphertext of a session-key $K \in \mathcal{K}$ under pk . We say that a KEM satisfies *correctness* if for all $k \in \mathbb{N}$, all keys (pk, sk) output from $\text{KKG}(1^k)$ and all ciphertext/session-key pairs (c, K) output from $\text{Encap}(pk)$, it holds that $\text{Decap}(sk, c) = K$.

Let $\text{ATK} \in \{\text{CPA}, 1\text{-CCA}, \text{CCA}\}$. We say that a KEM Γ is ATK secure if for all PPTAs \mathcal{A} , the advantage $\text{Adv}_{\Gamma, \mathcal{A}}^{\text{ATK}}(k) := 2 \cdot |\Pr[\text{Expt}_{\Gamma, \mathcal{A}}^{\text{ATK}}(k) = 1] - 1/2|$ is negligible, where the CCA experiment $\text{Expt}_{\Gamma, \mathcal{A}}^{\text{CCA}}(k)$ is defined as follows:

$$\begin{aligned} \text{Expt}_{\Gamma, \mathcal{A}}^{\text{CCA}}(k) : & [(pk, sk) \leftarrow \text{KKG}(1^k); (c^*, K_1^*) \leftarrow \text{Encap}(pk); K_0^* \leftarrow \{0, 1\}^k; \\ & b \leftarrow \{0, 1\}; b' \leftarrow \mathcal{A}^{\text{Decap}(sk, \cdot)}(pk, c^*, K_b^*); \text{Return } (b' \stackrel{?}{=} b)], \end{aligned}$$

where in the experiment, \mathcal{A} is not allowed to submit c^* to the oracle. The 1-CCA (1-bounded CCA) experiment $\text{Expt}_{\Gamma, \mathcal{A}}^{1\text{-CCA}}(k)$ is defined in the same way as the CCA

experiment, except that \mathcal{A} is allowed to submit a decapsulation query only once. Furthermore, the CPA experiment $\text{Expt}_{\mathcal{F}, \mathcal{A}}^{\text{CPA}}(k)$ is also defined similarly to the CCA experiment, except that \mathcal{A} is not allowed to submit any query.

Commitment. A commitment scheme \mathcal{C} consists of the two PPTAs (CKG, Com) with the following interface:

$$\begin{array}{ll} \textbf{Key Generation:} & \textbf{Commitment Generation:} \\ \hline ck \leftarrow \text{CKG}(1^k) & c \leftarrow \text{Com}(ck, m) \end{array}$$

where ck is a commitment key, and c is a commitment of the message m under ck .

As a (non-standard) requirement, we require the size of a commitment to be k -bit for k -bit security, no matter how long a committed message is.⁷

We say that a commitment scheme \mathcal{C} is *target-binding*⁸ if for all PPTAs $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, the advantage function $\text{Adv}_{\mathcal{C}, \mathcal{A}}^{\text{TBind}}(k) := \Pr[\text{Expt}_{\mathcal{C}, \mathcal{A}}^{\text{TBind}}(k) = 1]$ is negligible, where the experiment $\text{Expt}_{\mathcal{C}, \mathcal{A}}^{\text{TBind}}(k)$ is defined as follows:

$$\begin{array}{l} \text{Expt}_{\mathcal{C}, \mathcal{A}}^{\text{TBind}}(k) : [(m, r, \text{st}) \leftarrow \mathcal{A}_1(1^k); ck \leftarrow \text{CKG}(1^k); (m', r') \leftarrow \mathcal{A}_2(\text{st}, ck); \\ \text{Return } 1 \text{ iff } \text{Com}(ck, m'; r') = \text{Com}(ck, m; r) \wedge m' \neq m.] \end{array}$$

Since we do not directly use the hiding property, we do not introduce its formal definition. In Section 2.3, we define the trapdoor simulatability property for a commitment scheme, which is defined in essentially the same way as that for a TSPKE scheme.

⁷ This requirement (together with the following binding property and the “trapdoor simulatability property”) can be easily realized if we are given a TSPKE scheme and a UOWHF. We give the construction in the full version.

⁸ Note that the target-binding property is slightly weaker than the ordinary binding notion in the sense that an adversary has to choose its first message before seeing a key ck . The relation between the ordinary binding and target-binding is similar to the relation between collision resistance and target collision resistance of a hash function family. The target-binding was also used in [39].