Bounded CCA2-Secure Encryption

Ronald Cramer[1,2], Goichiro Hanaoka[3], Dennis Hofheinz[1], Hideki Imai[3,4], Eike Kiltz[1], Rafael Pass[5], abhi shelat[6], and Vinod Vaikuntanathan[7]

[1] Centrum voor Wiskunde en Informatica (CWI), Amsterdam
[2] Leiden University
[3] National Institute of Advanced Industrial Science and Technology, Tokyo
[4] Chuo University
[5] Cornell University
[6] University of Virginia
[7] Massachusetts Institute of Technology

Abstract. Whereas encryption schemes withstanding passive chosen-plaintext attacks (CPA) can be constructed based on a variety of computational assumptions, only a few assumptions are known to imply the existence of encryption schemes withstanding adaptive chosen-ciphertext attacks (CCA2). Towards addressing this asymmetry, we consider a weakening of the CCA2 model — *bounded CCA2-security* — wherein security needs only hold against adversaries that make an a-priori bounded number of queries to the decryption oracle. Regarding this notion we show (without any further assumptions):

- For any polynomial q, a simple *black-box* construction of q-bounded *IND-CCA2-secure* encryption schemes, from any *IND-CPA-secure* encryption scheme. When instantiated with the Decisional Diffie-Hellman (DDH) assumption, this construction additionally yields encryption schemes with very short ciphertexts.
- For any polynomial q, a (non-black box) construction of q-bounded *NM-CCA2-secure* encryption schemes, from any *IND-CPA-secure* encryption scheme. Bounded-CCA2 non-malleability is the strongest notion of security yet known to be achievable assuming only the existence of IND-CPA secure encryption schemes.

Finally, we show that non-malleability and indistinguishability are *not equivalent* under bounded-CCA2 attacks (in contrast to general CCA2 attacks).

1 Introduction

Encryption is often compared to a 'secure envelope'. Though appealing as a metaphor, understanding encryption requires a more formal *definition of security* of the primitive. For this task, the notion of *semantic security against adaptive chosen-ciphertext attacks* (in short, IND-CCA2 security) captures the essential characteristics of secure envelopes.

Under adaptive chosen-ciphertext attacks (CCA2), whose study was pioneered by Naor and Yung [22], and Rackoff and Simon [26], security is required

to hold with respect to adversaries that have access to a decryption oracle. This should be contrasted to the traditional type of chosen-plaintext attack (CPA), where the adversary is required to act on its own without any additional help [14].

While there are a number of candidate (practical) public-key encryption schemes known to be semantically secure against a CPA attack [13], designing ones that withstand a CCA2 attack is a delicate and difficult task. In the standard model, there are essentially three approaches known. The first approach, put forth by Naor and Yung [22] in the early 1990s, and subsequently extended by Dolev, Dwork and Naor [10], and later Sahai [28] and Lindell [20], is based on the use of non-interactive zero knowledge for NP. This leads to schemes based on quite general cryptographic assumptions. The second is due to Cramer and Shoup [6–8] and is based on hash-proof systems. This leads to quite practical schemes based on several concrete number-theoretic assumptions. The third and most recent approach is due to Canetti, Halevi and Katz [3] and relies on identity-based cryptography.

To sum up, all the above approaches make use of additional assumptions to construct CCA2-secure schemes (apart from the existence of CPA-secure encryption schemes). A fundamental open question is thus:

> *Can any CPA-secure encryption scheme be transformed into one that is also CCA2 secure, without making additional complexity-theoretic assumptions?*

1.1 Our Results

Towards addressing this fundamental question, in this paper we introduce a weakening of the CCA2 attack which we call a *bounded-CCA2* attack. In such an attack, the adversary is restricted to making an *a-priori bounded* number of queries to the decryption oracle. This is indeed a reasonable model, since the use of encryption in many protocols (such as secure multiparty computation) can be upper-bounded to q decryptions. With this terminology, our main contributions are summarized below. Henceforth, unless otherwise mentioned, whenever we talk of CCA attacks, we mean adaptive chosen ciphertext attacks (CCA2), as opposed to the weaker lunch-time attacks (CCA1).

Bounded CCA2 Semantic Security. Our first result is a simple and efficient *black-box* construction of a public-key encryption (PKE) scheme that is semantically secure against a q-bounded CCA2 attack (technically termed IND-q-CCA-secure), starting from any CPA-secure encryption scheme. Technically, this result combines techniques from [3, 9]. However, it appears that the implications for black-box constructions of chosen ciphertext secure encryption from semantically secure encryption, as we deduce them here, have not been reported before.

Theorem 1 (Informal). *For any polynomial q, there exists a black-box construction of an* IND-q-CCA*-secure encryption scheme from any CPA-secure encryption scheme.*

The key size and the ciphertext size of this construction are quadratic in q and thus quite large; nevertheless, it demonstrates the *feasibility* of black-box constructions of bounded-CCA2-secure encryption schemes from any CPA-secure scheme. Interestingly, this result stands in sharp contrast to the recent results of Gertner, Malkin and Myers [12] showing that "such" black-box constructions are impossible when considering standard (unbounded) CCA2-secure encryption. (The black-box separation result from [12] only holds for constructions where the *decryption* function of the CCA2 secure scheme does not make calls to the *encryption* function of the CPA secure scheme. Our black-box construction of q-bounded CCA2 secure encryption falls into this category.)

We additionally show that if the underlying CPA-secure PKE scheme has certain homomorphic properties, then we can construct a q-bounded CCA2-secure PKE scheme with very short ciphertexts. In particular, in groups where the DDH assumption holds, we can give a q-bounded CCA2 secure PKE scheme with only one group element of ciphertext expansion. In contrast, the best known DDH-based schemes such as the one by Kurosawa and Desmedt [18] which achieve full CCA2 security have two group elements plus a MAC. The length of the public keys in this construction are, however, still quadratic in q.

Bounded CCA2 Non-malleability. A q-bounded-CCA2 non-malleable (in technical terms, NME-q-CCA-secure) encryption scheme is one that is "non-malleable" with respect to an adversary making at most q decryption queries. For this notion, we are able to show:

Theorem 2 (Informal). *Assuming CPA-secure public-key encryption schemes exist, for any polynomial q, there exists an* NME-q-CCA*-secure encryption scheme.*

As far as we know, the notion of bounded-CCA2 non-malleability is the strongest notion of security for encryption schemes known to be achievable under only the assumption of CPA-secure encryption schemes. Furthermore, the length of both the the public-key and the ciphertexts grows *linearly* with q (instead of quadratically as in our previous construction). However, this second construction makes a *non-black-box* use of the underlying CPA secure encryption scheme. In particular, we use a proof that several ciphertexts are encryptions of the same message, and this may require analyzing the encryption circuit to form a theorem statement. (On the other hand, even though our construction uses ZK proofs and thus costly $\mathcal{NP}$ reductions, in many cases, there exist efficient proofs — Σ protocols [4], for example — for the type of theorems we encounter.)

Relation Between Semantic Security and Non-malleability Against Bounded CCA2 Attacks. It is known that under a CCA2 attack, the otherwise weaker notion of semantic security in fact implies also non-malleability [1]. In the case of bounded-CCA2 security, however, we show that this equivalence does not hold. In particular, we show that q-bounded-CCA2 security for any (fixed) q does not even imply non-malleability under the simple CPA attack.

Theorem 3 (Informal). *Assume CPA-secure public-key encryption schemes exist. Then, for every q, there exists an encryption scheme that is q-bounded CCA2-secure, but is not non-malleable (even under a CPA attack).*

This separation of notions highlights the importance of directly proving non-malleability of our second scheme (which slightly complicates the analysis).

1.2 Importance of These Results

The notion of *bounded* CCA2 security which we present is a weakening of the traditional notion of CCA2 security. Since it is possible to achieve CCA2 security, one may then wonder why it is important to consider this notion. There are in fact two simple reasons:

1. There are many hardness assumptions (such as computational-Diffie-Hellman and many lattice-based hardness assumptions) for which we can only construct CPA-secure encryption schemes. Our results show how to transform those schemes into ones with much stronger security properties. Since no one knows how to achieve full (unbounded) CCA2 security under these assumptions, our result represents the state-of-the-art for encryption in that area.
2. Being a weaker notion, bounded-CCA2 security may allow for more efficient constructions. Indeed, under the DDH assumption, we present a bounded-CCA2 scheme which is less than half the size of the smallest full-CCA2 secure scheme. For certain low-bandwidth applications in which the size of the ciphertext is critical, this may be the best construction to use.

Organization. After fixing some notation in §2, we formally define the notion of q-bounded CCA2 security. Section §3 contains a black-box construction of a q-bounded IND-CCA-secure encryption scheme, and Section §4 contains an optimized instantiation under the DDH assumption. Section §5 contains a non-black-box construction of a q-bounded NME-CCA-secure encryption scheme. Finally, in Section §6, we present a separation between the definitions of semantic security and non-malleability under q-bounded attacks.

Publication Info. This paper is a merge of three independent preprints [5, 15, 23].

2 Preliminaries and Definitions

If S is a set then $s \xleftarrow{\$} S$ denotes the operation of picking an element s of S uniformly at random. We write $\mathcal{A}(x, y, \ldots)$ to indicate that $\mathcal{A}$ is an algorithm with inputs $x, y, \ldots$ and by $z \xleftarrow{\$} \mathcal{A}(x, y, \ldots)$ we denote the operation of running $\mathcal{A}$ with inputs $(x, y, \ldots)$ and letting z be the output. We write $\mathcal{A}^{\mathcal{O}_1, \mathcal{O}_2, \cdots}(x, y, \ldots)$ to indicate that $\mathcal{A}$ is an algorithm with inputs $x, y, \ldots$ and black-box access to oracles $\mathcal{O}_1, \mathcal{O}_2, \ldots$. If $\mathcal{A}$ is a randomized algorithm, the notation $\mathcal{A}(x; r)$ means running $\mathcal{A}$ with input x and randomness r.

Definition 1 (Encryption scheme). *A triple* PKE $= (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is a public key encryption scheme, if (1)* Gen *and* Enc *are p.p.t. algorithms and*

Dec *is a deterministic polynomial-time algorithm, (2)* Gen *on input* 1^k *produces a pair* (pk, sk)*, where* pk *is the public-key and* sk *is the secret-key, (3)* $\mathsf{Enc} : pk \times \{0,1\}^* \to \{0,1\}^*$ *runs on input a public key* pk *and a message* $m \in \{0,1\}^*$ *and produces a ciphertext* c*, (4)* $\mathsf{Dec} : sk \times \{0,1\}^* \to \{0,1\}^* \cup \{\perp\}$ *runs on input* (sk, c) *and produces either a message* $m \in \{0,1\}^*$ *or a special symbol* $\perp$*, (5)*(Perfect Correctness) *There exists a polynomial* $p(k)$ *and a negligible function* $\mu(k)$ *such that for* every message m*, and* every random tape r_e*,*

$$\Pr[(pk, sk) \leftarrow \mathsf{Gen}(1^k; r_g) : \ \exists r_e, m \ s.t \ \mathsf{Dec}_{sk}(\mathsf{Enc}_{pk}(m; r_e)) \neq m] \leq \mu(k).$$

where the probability is over the random choice of r_g*. That is, with high probability over the keys generated by* Gen*, all valid ciphertexts decrypt correctly.*

Next, we define the notions of IND-q-CCA-security and NME-q-CCA-security.

Definition 2 (IND-q-CCA security). *For a function* $q(k) : \mathbb{N} \to \mathbb{N}$*, we define the security notion of indistinguishability against* q*-bounded CCA adversaries* (IND-q-CCA)*. For an adversary* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ *we define the advantage function*

$$\mathbf{Adv}^{\mathsf{IND\text{-}q\text{-}CCA}}_{\mathrm{PKE},\mathcal{A}}(k) \ = \ \left| \Pr[\mathbf{Exp}^{\mathsf{IND}\text{-}q\text{-}\mathsf{CCA}\text{-}1}_{\mathrm{PKE},\mathcal{A}}(k) = 1] - \Pr[\mathbf{Exp}^{\mathsf{IND}\text{-}q\text{-}\mathsf{CCA}\text{-}0}_{\mathrm{PKE},\mathcal{A}}(k) = 1] \right|$$

where, for $b \in \{0,1\}$*,* $\mathbf{Exp}^{\mathsf{IND}\text{-}q\text{-}\mathsf{CCA}\text{-}b}_{\mathrm{PKE},\mathcal{A}}$ *is defined by the following experiment.*

Experiment $\mathbf{Exp}^{\mathsf{IND}\text{-}q\text{-}\mathsf{CCA}\text{-}b}_{\mathrm{PKE},\mathcal{A}}(k)$
$(pk, sk) \xleftarrow{\$} \mathsf{Gen}(1^k)$
$(M_0, M_1, St_1) \xleftarrow{\$} \mathcal{A}_1^{\mathsf{Dec}(sk,\cdot)}(pk)$ *s.t.* $|M_0| = |M_1|$
$c^* \xleftarrow{\$} \mathsf{Enc}(pk, M_b)$
$b' \xleftarrow{\$} \mathcal{A}_2^{\mathsf{Dec}(sk,\cdot)}(c^*, St_1)$
Return b'

The adversary $(\mathcal{A}_1, \mathcal{A}_2)$ *is restricted to ask at most* $q(k)$ *queries to the decryption oracle* Dec *in total in each run of the experiment, and none of the queries of* $\mathcal{A}_2$ *may contain* c^**. The scheme* PKE *is said to be indistinguishable against* q*-bounded chosen-ciphertext attacks (*IND-q-CCA*-secure, in short) if the advantage function* $\mathbf{Adv}^{\mathsf{IND\text{-}q\text{-}CCA}}_{\mathrm{PKE},\mathcal{A}}(k)$ *is negligible in* k *for all adversaries* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*.*

We have the following relation to the standard security definitions for PKE schemes. Scheme PKE is said to be (1) indistinguishable against chosen-plaintext attacks [14] (CPA), denoted IND-CPA, if it is IND-0-CCA-secure, and (2) indistinguishable against chosen-ciphertext attacks [26] (CCA2), denoted IND-CCA, if it is IND-q-CCA-secure for *any* polynomial $q(k)$.

As was done above with indistinguishability, we extend the definition of non-malleability presented in [24] to consider $q(k)$-bounded adversaries.

Definition 3 (NME-q-CCA security). *Let* PKE $= (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be an encryption scheme and let the random variable* $\mathsf{NME\text{-}q\text{-}CCA_b}(\Pi, \mathsf{A}, \mathsf{k}, \ell)$ *where* $b \in$

$\{0,1\}$, $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ *and* $k, \ell \in \mathbb{N}$ *denote the result of the following probabilistic experiment:*

$\mathsf{NME\text{-}q\text{-}CCA_b}(\mathrm{PKE}, \mathcal{A}, \mathsf{k}, \ell)$:

$(pk, sk) \leftarrow \mathsf{Gen}(1^k)$

$(m_0, m_1, \mathrm{STATE}_{\mathcal{A}}) \leftarrow \mathcal{A}_1^{\mathsf{Dec}(sk,\cdot)}(pk)$ *s.t.* $|m_0| = |m_1|$

$y \leftarrow \mathsf{Enc}_{pk}(m_b)$

$(c_1, \ldots, c_\ell) \leftarrow \mathcal{A}_2^{\mathsf{Dec}(sk,\cdot)}(y, \mathrm{STATE}_{\mathcal{A}})$

Output $(d_1, \ldots, d_\ell)$ *where* $d_i = \begin{cases} \mathrm{COPY} & \text{if } c_i = y \\ \mathsf{Dec}_{sk}(c_i) & \text{otherwise} \end{cases}$

$\mathrm{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is* $\mathsf{NME\text{-}q\text{-}CCA}$*-secure for a function* $q(k) : \mathbb{N} \to \mathbb{N}$ *if,* $\forall$ *p.p.t. algorithms* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ *which make* $q(k)$ *total queries to the oracles and for any polynomial* $p(k)$*, the following two ensembles are computationally indistinguishable:*

$$\left\{\mathsf{NME\text{-}q\text{-}CCA_0}(\mathrm{PKE}, \mathcal{A}, \mathsf{k}, \mathsf{p}(\mathsf{k}))\right\}_{\mathsf{k}\in\mathbb{N}} \overset{\mathsf{c}}{\approx} \left\{\mathsf{NME\text{-}q\text{-}CCA_1}(\mathrm{PKE}, \mathcal{A}, \mathsf{k}, \mathsf{p}(\mathsf{k}))\right\}_{\mathsf{k}\in\mathbb{N}}$$

If $q(k) = 0$, then the encryption scheme is said to be $\mathsf{NME\text{-}CPA}$-secure.

3 Construction of Bounded IND-CCA Secure Encryption

In this section, we present a black-box construction of an $\mathsf{IND\text{-}q\text{-}CCA}$-secure encryption scheme. The general outline of our construction is as follows.

First, as demonstrated by Canetti, Halevi and Katz [3], every identity-based encryption scheme can be transformed into a fully chosen-ciphertext secure encryption scheme. Second, an $\mathsf{IND\text{-}CPA}$ secure encryption scheme implies a "q-resilient" identity-based encryption scheme. (The notion of q-resilient security in the context of identity-based encryption [16] means that the scheme guarantees security as long as at most q private keys are established.) The latter result is only implicitly contained in a paper about key-insulated public-key cryptosystems by Dodis, Katz, Xu, and Yung [9]. A closer observation of the combination of the two results already reveals the construction of our $\mathsf{IND\text{-}q\text{-}CCA}$-secure encryption scheme. Since both transformations are black-box, our main result can be obtained. However, it appears that the implications for black-box constructions of $\mathsf{IND\text{-}q\text{-}CCA}$-secure encryption from $\mathsf{IND\text{-}CPA}$-secure encryption as we deduce them here have not been reported before.

Stateful versus Stateless Encryption. When one only considers stateful encryption, the problem of constructing black-box $\mathsf{IND\text{-}q\text{-}CCA}$-secure encryption becomes trivial: the receiver's public-key contains q independent public-keys pk_i of the $\mathsf{IND\text{-}CPA}$-secure scheme. For $1 \leq j \leq q$, to encrypt the j^{th} message, a sender uses the j^{th} public-key pk_j as a "one-time key" for the $\mathsf{IND\text{-}CPA}$-secure encryption scheme, the state being j that is incremented after each encryption. However, this construction requires *all participants* to share and update the dynamic state information j. (This is in contrast to signature schemes where the signer may maintain a private state.)

We circumvent this unrealistic state update assumption by "load-balancing" the use of instances of the IND-CPA-secure base scheme. The general outline of our construction is as follows. We use the q-resilient identity-based encryption construction implicitly given in [9] based on any IND-CPA-secure PKE scheme. Using a transformation from [3], this q-resilient identity-based encryption scheme can be transformed into a PKE scheme. As we will see, the resulting PKE scheme is secure against q-bounded chosen-ciphertext adversaries.

Theorem 4. *For any fixed polynomial q, there exists a black-box construction that, given any* IND-CPA*-secure scheme* $(\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$*, builds an* IND-q-CCA*-secure public-key encryption scheme* $(\mathsf{Gen}^{\mathsf{kg}}, \mathsf{Enc}^{\mathsf{kg},\mathsf{enc}}, \mathsf{Dec}^{\mathsf{kg},\mathsf{dec}})$.

Here we give a direct proof of this theorem that bypasses the notion of identity-based encryption altogether. We furthermore note that there are some technical problems with the security proof of the implicitly contained q-resilient IBE scheme from [9] that we fix in this note.[8]

3.1 Building Blocks

Cover-free families. If S, T are sets, we say that S does not cover T if $S \not\supseteq T$. Let d, q, s be positive integers, and let $F = (F_i)_{1 \le i \le s}$ be a family of subsets of $\{1, \ldots, d\}$. We say that family F is *q-cover-free over* $\{1, \ldots, d\}$, if for each subset $F_i \in F$ and each S that is the union of at most q sets in $(F_1, \ldots, F_{i-1}, F_{i+1}, \ldots, F_s)$, it is the case that S does not cover F_i. Furthermore, we say that F is l-uniform if all subsets in the family have size l. We use the following fact [11, 17]: there is a deterministic polynomial time algorithm that on input integers s, q returns l, d, F where $F = (F_i)_{1 \le i \le s}$ is a l-uniform q-cover-free family over $\{1, \ldots, d\}$, for $l = d/4q$ and $d \le 16q^2 \log(s)$. In the following we let SUB denote the resulting deterministic polynomial-time algorithm that on input s, q, i returns F_i. We call $F_i = \mathsf{SUB}(s(k), q(k), i)$ the subset associated to index $i \in \{1, \ldots, s(k)\}$.

For our construction we will need a cover-free family with the parameters

$$s(k) = 2^k, \quad d(k) = 16kq^2(k), \quad l(k) = 4kq(k)\ . \tag{1}$$

One-time signatures. In our construction, we need a strong one-time signature scheme $\text{OTS} = (\mathsf{Sigkg}, \mathsf{Sign}, \mathsf{Verify})$ (see [19]). We assume that the verification keys which are part of the output by Sigkg are bit strings of size k which we interpret as natural numbers in $\{1, \ldots, 2^k\}$. Strong one-time signature schemes can be constructed from (the key-generation algorithm of) any IND-CPA-secure

[8] The problem in the proof of Theorem 2 in [9] (only contained in the full version) is that their simulator (simulating the view of an adversary attacking the IBE scheme) sometimes is forced to abort. However, this forced abort is *n*ot independent of the adversary's view in this simulation. This dependence could be exploited by an adversary that has a higher chance in breaking the IBE scheme only if the simulator aborts. We give a different simulation to overcome this problem.

encryption scheme via a black-box reduction (since a one-way function can be constructed from the key-generation algorithm, and one-way functions imply strong signature schemes [19, 27]).

3.2 The Construction

Let $q(k) : \mathbb{N} \to \mathbb{N}$ be a function. Our construction of the IND-q-CCA encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with black-box access to the IND-CPA-secure encryption scheme $(\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ is depicted in Fig. 1. In general we can also use any computationally secure all-or-nothing transform (e.g., the black-box construction from [2] based on one-way functions) to decrease ciphertext size.

Public and secret keys have size polynomial (quadratic) in the maximal number of decryption queries $q(k)$. Also note that the upper bound $q(k)$ must be known in advance as a parameter of the construction.

$\mathsf{Gen}^{\mathsf{kg}}(1^k)$: Define $s(k) = 2^k, d(k) = 16kq^2(k), l(k) = 4kq(k)$ as in Equation (1). For $i = 1, \ldots, d(k)$ run $(pk_i, sk_i) \xleftarrow{\$} \mathsf{kg}(1^k)$. Output $PK = (pk_1, \ldots, pk_{d(k)})$ and $SK = (sk_1, \ldots, sk_{d(k)})$.

$\mathsf{Enc}^{\mathsf{kg},\mathsf{enc}}(PK, M)$**:** Create a random pair of one-time signing keys $(vk, sigsk) \xleftarrow{\$} \mathsf{Sigkg}^{\mathsf{kg}}(1^k)$. Let $F_{vk} = \{s_1, \ldots, s_{l(k)}\}$ be the subset associated to verification key vk. Pick random $M_1, \ldots, M_{l(k)}$ subject to $M = M_1 \oplus \ldots \oplus M_{l(k)}$ and run $c_j \xleftarrow{\$} \mathsf{enc}(pk_{s_j}, M_j)$, for $j = 1, \ldots, l(k)$. Sign the ciphertexts $c = (c_1, \ldots, c_{l(k)})$ with $sigsk$ by running $\sigma \leftarrow \mathsf{Sign}^{\mathsf{kg}}(sigsk, c)$ and output $C = (c, vk, \sigma)$.

$\mathsf{Dec}^{\mathsf{kg},\mathsf{dec}}(SK, (c = (c_1, \ldots, c_{l(k)}), vk, \sigma))$**:** If $\mathsf{Verify}^{\mathsf{kg}}(vk, c, \sigma)$ rejects, return *reject*. Let $F_{vk} = \{s_1, \ldots, s_{l(k)}\}$ be the subset associated to vk. For $j = 1, \ldots, l(k)$ run $M_j \leftarrow \mathsf{dec}(sk_{s_j}, c_j)$ and output $M = M_1 \oplus \ldots \oplus M_{l(k)}$.

Fig. 1. Black-box construction of an IND-q-CCA secure encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ from any IND-CPA-secure scheme $(\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$

The following proves our main result, Theorem 4.

Lemma 1. *If* $(\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ *is* IND-CPA *secure then* $(\mathsf{Gen}^{\mathsf{kg}}, \mathsf{Enc}^{\mathsf{kg},\mathsf{enc}}, \mathsf{Dec}^{\mathsf{kg},\mathsf{dec}})$ *as described in Fig. 1 is* IND-q-CCA *secure.*

Proof. For any PPT adversary $\mathcal{A}$ against the IND-q-CCA security of $(\mathsf{Gen}^{\mathsf{kg}}, \mathsf{Enc}^{\mathsf{kg},\mathsf{enc}}, \mathsf{Dec}^{\mathsf{kg},\mathsf{dec}})$, we show, via a game-based proof, that $\mathcal{A}$'s advantage in the IND-q-CCA game is negligible.

Let **Game 0** be the IND-q-CCA game with adversary $\mathcal{A}$ and uniformly chosen experiment bit b. Let X_0 denote the event that $\mathcal{A}$'s final guess is correct (i.e., X_0 denotes that $b' = b$). For later games, let X_i $(i > 0)$ be defined analogously.

$$\frac{1}{2}\mathbf{Adv}^{\mathsf{IND\text{-}q\text{-}CCA}}_{\mathrm{PKE},\mathcal{A}}(k) = |\Pr[X_0] - \frac{1}{2}| \,.$$

Game 1 is identical to **Game 0**, except that the verification key vk^* for the challenge ciphertext is *initially* chosen, and all decryption queries with $vk = vk^*$ are rejected.

By reduction on the security of the signature scheme OTS, one can show that

$$|\Pr[X_1] - \Pr[X_0]| \leq \mathbf{Adv}^{ots\text{-ex-for}}_{\mathrm{OTS},\mathcal{F}}(k) ,$$

for a suitable adversary $\mathcal{F}$, where $\mathbf{Adv}^{ots\text{-ex-for}}_{\mathrm{OTS},\mathcal{F}}(k)$ is the probability that $\mathcal{F}$ breaks the existential unforgeability of the one-time signature scheme.

Game 2 proceeds like **Game 1**, but we introduce some notation useful for later. Denote by $C^{(i)} = (c^{(i)}, vk^{(i)}, \sigma^{(i)})$ the i-th decryption request of $\mathcal{A}$. Define

$$Q := \bigcup_{vk^{(i)} \neq vk^*} F_{vk^{(i)}}$$

for the sets $F_{vk^{(i)}}$ of PKE keypairs associated with the respective i-th query. We know that $F_{vk^*} \not\subseteq Q$, so we can define $j := \min(F_{vk^*} \setminus Q)$. Additionally, we choose (this can be done at the beginning of the game, right after vk^* is fixed) uniformly and independently $i \in F_{vk^*}$. Call FAIL the event that $i \neq j$. Note that

$$\Pr[\mathrm{FAIL} \mid X_2] = \frac{l-1}{l} = \Pr[\mathrm{FAIL}],$$

so the events X_2 and FAIL are independent, and in particular, $\Pr[X_2] = \Pr[X_2 \mid \neg\mathrm{FAIL}]$. Since we did not actually change anything, $\Pr[X_2] = \Pr[X_1]$.

In **Game 3**, we substitute $\mathcal{A}$'s output b' with a random bit whenever FAIL occurs. Obviously,

$$\Pr[X_3 \mid \neg\mathrm{FAIL}] = \Pr[X_2 \mid \neg\mathrm{FAIL}] \text{ and } \Pr[X_3 \mid \mathrm{FAIL}] = \frac{1}{2}.$$

Since $\Pr[\mathrm{FAIL}] = (l-1)/l$ in Game 3 as well, we can establish that

$$\Pr[X_3] - \frac{1}{2} = \frac{\Pr[X_2] - \frac{1}{2}}{l}.$$

In **Game 4**, we immediately stop the experiment and set FAIL to true (hence immediately taking a random bit for $\mathcal{A}$'s output) as soon as $\mathcal{A}$ asks for a decryption of a ciphertext with a verification key $vk \neq vk^*$ such that $i \in F_{vk}$. Note that already in Game 3, such a query would have implied $j \neq i$ and hence FAIL. Consequently,

$$\Pr[X_4] = \Pr[X_3].$$

Note that Game 4 can be run without knowledge of the secret key sk_i.

In **Game 5**, the challenge ciphertext is prepared as follows. For encrypting the challenge message M_b with PKE, we first choose uniformly PKE plaintexts $M_1^*, \ldots, M_{i-1}^*, M_{i+1}^*, \ldots, M_l^*$ and *then* the suitable

$$M_i^* := M_b \oplus \bigoplus_{r \neq i} M_r^*.$$

Note that then, only the plaintext M_i^* depends on the experiment bit b. This does not change the distribution of the whole vector $M_1^*, \ldots, M_l^*$, and we have

$$\Pr[X_5] = \Pr[X_4].$$

On the other hand, Game 5 can be directly mapped to an adversary $\mathcal{B}$ on the IND-CPA security of PKE. More concretely, $\mathcal{B}$ simulates Game 5, but substitutes pk_i with its own challenge public key, and submits as challenge plaintexts

$$\hat{M}_0 := M_0 \oplus \bigoplus_{r \neq i} M_r^* \text{ and } \hat{M}_1 := M_1 \oplus \bigoplus_{r \neq i} M_r^*.$$

Then, $\Pr[X_5]$ is precisely the success probability of $\mathcal{B}$ in the IND-CPA experiment

$$|\Pr[X_5] - \frac{1}{2}| = \frac{1}{2}\mathbf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathrm{PKE},\mathcal{B}}(k) .$$

Collecting probabilities shows that

$$\mathbf{Adv}^{\mathsf{IND\text{-}q\text{-}CCA}}_{\mathrm{PKE},\mathcal{A}}(k) \leq l(k) \cdot \mathbf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathrm{PKE},\mathcal{B}}(k) + 2 \cdot \mathbf{Adv}^{ots\text{-}\mathrm{ex}\text{-}\mathrm{for}}_{\mathrm{OTS},\mathcal{F}}(k) .$$

Since $\mathbf{Adv}^{\mathsf{IND\text{-}CPA}}_{\mathrm{PKE},\mathcal{B}}$ and $\mathbf{Adv}^{ots\text{-}\mathrm{ex}\text{-}\mathrm{for}}_{\mathrm{OTS},\mathcal{F}}$ are negligible, this shows the claim. □

Remark 1. We stress that it is important for our construction that the number of subsets $s(k)$ is super-polynomial in k. One could try to trivially build $q(k)$-bounded CCA secure encryption PKE from CPA secure PKE using a public/secret key vector of size $q(k)$ and defining the subsets F_i as $\{i\}$, for $1 \leq i \leq s(k) := q(k)$. For encryption, a message gets encrypted using pk_{vk}, where $vk \in \{1, \ldots, q(k)\}$ is one of the $q(k)$ distinct public keys of PKE, and vk is a random verification key of the signature scheme. However, since there are only $q(k)$ many possible choices of verification keys, one can break the scheme with probability $\frac{1}{q(k)}$ by (trivially) breaking the signature scheme with probability $\frac{1}{q(k)}$.

Remark 2. It might be interesting to explore what (additional) security properties PKE satisfies once invoked with a scheme PKE that itself is not only IND-CPA-secure, but, say, NME-CPA-secure. Unfortunately, we cannot hope that PKE is NME-CPA-secure, independently of PKE's security: say that adversary $\mathcal{A}$ receives a challenge ciphertext $C^* = (c^*, vk^*, \sigma^*)$ with $c^* = (c_1, \ldots, c_l)$ and $F_{vk^*} = \{s_1^*, \ldots, s_l^*\}$. Then $\mathcal{A}$ may be able to construct $l(k)$ ciphertexts $C^{(1)}, \ldots, C^{(l)}$ such that $C^{(i)}$ is associated with a subset $F^{(i)}$ with $s_i^* \in F^{(i)} \neq F_{\mathrm{verk}}$, and the vector $c^{(i)}$ consists only of 0-encryptions except for c_i^*. The XOR of the decryptions of $C^{(i)}$ is precisely the challenge plaintext, hence this is a successful malleability attack.

We note that if we assume the IND-CCA1 security of PKE, this proof also shows that the resulting scheme PKE is secure against IND-CCA attackers who have full access to a decryption oracle before receiving the challenge ciphertext, but only limited access (q queries) to it in the second attack phase.

4 Bounded IND-CCA-secure Encryption from DDH

In this section we propose a construction of IND-q-CCA-secure encryption based on the Decisional Diffie Hellman (DDH) assumption. The construction follows the approach from the previous section; we make use of cover-free sets and (with the same parameters as in Section 3) set up $d(k)$ independent instances of the (semantically secure) El-Gamal encryption scheme. We encrypt a message using a subset of the $d(k)$ keys, where the subset is determined by cover-free sets. Certain homomorphic properties of El-Gamal encryption are exploited to shrink the ciphertext size down to one group element. (This stands in contrast to Cramer-Shoup encryption which requires 4 group elements, and the Kurosawa-Desmedt one which requires 2 group elements and a MAC.) The main contribution of this section is to demonstrate the existence of such limited $q(k)$-bounded CCA secure schemes with such an optimal ciphertext size.
To instantiate our scheme we need the following building blocks:

- A cyclic group $\mathbb{G}$ of prime-order p where the DDH assumption is believed to hold, i.e, the two distributions (g, g^x, g^y, g^{xy}) and (g, g^x, g^y, g^z) are computationally indistinguishable, for random $g \in \mathbb{G}$, and random $x, y, z \in \mathbb{Z}_p$.
- A redundancy-free symmetric-key encryption scheme (E, D) which is secure against chosen-ciphertext attacks [8]. Such schemes can be constructed based on strong pseudorandom permutations [25]. For simplicity, we assume that the key space of (E, D) is $\mathbb{G}$. (In practice, we can convert $K \in \mathbb{G}$ into a random binary string by using key derivation functions [8].)
- A hash function $\mathsf{TCR} : \mathbb{G} \to \{0,1\}^k$ that is assumed to be target collision-resistant [21].

Let $\mathbb{G}$ be a prime order group and g a random generator of $\mathbb{G}$. The construction is given in Fig. 2. Correctness is easy to verify. Public and secret keys have quadratic size in the maximal number of decryption queries $q(k)$. The ciphertext overhead of the scheme (i.e., the difference between ciphertext and plaintext size) is only one group element $c \in \mathbb{G}$. The ciphertext length of our scheme is considered optimal since it is the same as that of the CPA secure (original) El-Gamal encryption.

Theorem 5. *Assume* TCR *is a target collision-resistant hash function,* $\mathbb{G}$ *is a group where the DDH assumption holds, and* (E, D) *is a symmetric encryption scheme that is secure against chosen-ciphertext attacks. Then* PKE *as described in Fig. 2 satisfies* IND-q-CCA *security.*

The proof of this theorem is very similar to the one of Lemma 1 and is omitted here. The idea is to prove that the underlying key encapsulation mechanism (KEM) is IND-q-CCA-secure under the DDH assumption. Using the KEM/DEM composition theorem [8], this implies the result. Intuitively, we can explain $q(k)$-bounded CCA security of the KEM part as follows: Given $(g, g^x, g^y, h) \in \mathbb{G}^4$, an algorithm $\mathcal{B}$ against the DDH problem randomly picks α from F_{t^*} where $t^* = \mathsf{TCR}(g^y)$, and sets $X_\alpha \leftarrow g^x$. For all $i \in \{1, \ldots, d(k)\} \backslash \{\alpha\}$, $\mathcal{B}$ computes

$\mathsf{Gen}(1^k)$: Define $s(k) = 2^k, d(k) = 16kq^2(k), l(k) = 4kq(k)$. For $i = 1, \ldots, d(k)$ compute $X_i = g^{x_i}$ for random $x_i \in \mathbb{Z}_p$. Output $PK = (X_1, \ldots, X_{d(k)})$ and $SK = (x_1, \ldots, x_{d(k)})$.

$\mathsf{Enc}(PK, M)$**:** Compute $c = g^r$ for random $r \in \mathbb{Z}_p$. Let F_t be the subset associated to $t = \mathsf{TCR}(c)$. Use symmetric key $K = (\prod_{i \in F_t} X_i)^r$ to encrypt message M to $\psi \leftarrow \mathsf{E}_K(M)$. Output $C = (c, \psi)$.

$\mathsf{Dec}(SK, C = (c, \psi))$**:** Let F_t be the subset associated to $t = \mathsf{TCR}(c)$. Reconstruct the symmetric key as $K = c^{\sum_{i \in F_t} x_i}$ and decrypt ψ to $M \leftarrow \mathsf{D}_K(\psi)$.

Fig. 2. An IND-q-CCA-secure PKE scheme based on DDH.

$x_i \xleftarrow{\$} \mathbb{Z}_p^*$ and $X_i \leftarrow g^{x_i}$, and gives $PK = (X_1, \ldots, X_{d(k)})$ to another adversary $\mathcal{A}$ against the IND-q-CCA security of the KEM part. $\mathcal{B}$ also sets (c^*, K^*) as a challenge which will be given to $\mathcal{A}$, where $c^* = g^y$, and $K^* = h \cdot \prod_{i \in F_{t^*} \setminus \{\alpha\}} (g^y)^{x_i}$. $\mathcal{B}$ outputs "$h = g^{xy}$" if $\mathcal{A}$ outputs "real key", or "$h \neq g^{xy}$" otherwise. It is clear that for any query c, $\mathcal{B}$ can respond $K = c^{\sum_{i \in F_t} x_i}$ unless $\alpha \in F_t$ where $t = \mathsf{TCR}(c)$. Then, by a similar argument to that in Lemma 1, we can show that $\mathcal{B}$ breaks the DDH assumption.

5 Construction of Bounded NME-CCA-Secure Encryption

In this section, we construct an NME-q-CCA-secure encryption scheme using any semantically secure (IND-CPA-secure) encryption scheme. The construction is *the same* as the DDN construction [10] and the construction of Pass, Shelat and Vaikuntanathan [24], except that the NIZK proof used is a "designated-verifier" NIZK proof (DV-NIZK) with "q-bounded strong soundness". Informally, a designated-verifier NIZK proof is one where the verifier has some secret information that enables him to check the validity of a proof. A DV-NIZK proof is q-bounded sound, if soundness holds even against an adversary who can query the verifier on at most q proofs and learn if the proofs are valid or not. We refer the reader to the full version for definitions and constructions of such designated verifier NIZK (relying on the construction from [24]).[9]

Because the security proof for this construction is so similar to the one from [24], we merely summarize the differences necessary to take care of the additional decryption oracle available to a q-CCA adversary. For a full proof, refer to the full version of this paper.

Theorem 6. *Assume there exists an* IND-CPA*-secure scheme. Then, for every polynomial q, there exists an encryption scheme that is* NME-q-CCA*-secure.*

[9] For technical reasons we also require to slightly strengthen the zero-knowledge property of designated verifier NIZK of [24].

Proof idea: Recall that an encryption of a message m from the construction in [24] is of the form $(\mathbf{c}, \pi, vk, \sigma)$, where $vk := v_1 \dots v_k$ is a k-bit verification-key for a strong one-time signature scheme, $\mathbf{c} = (c_1, \dots, c_k)$ is a vector of encryptions of m where c_i is an encryption of m under the $\mathsf{IND\text{-}CPA}$ public-key pk_{v_i}, π is a DV-NIZK proof that all the encryptions in $\mathbf{c}$ are encryptions of the same message, and σ is a signature of $(\mathbf{c}, \pi)$ under a signing key corresponding to vk.

The proof in [24] proceeds by defining hybrid experiments $\mathsf{NME}_b^{(1)}$ and $\mathsf{NME}_b^{(2)}$ and proceeding to show that the experiments are indistinguishable, and that if an adversary succeeds in breaking $\mathsf{NME}_b^{(2)}$, it breaks the semantic security of the underlying encryption scheme.

We will proceed in a completely analogous way, by defining experiments $\mathsf{NME\text{-}q\text{-}CCA}_\mathsf{b}^{(1)}$ and $\mathsf{NME\text{-}q\text{-}CCA}_\mathsf{b}^{(2)}$ for $b \in \{0, 1\}$. The experiment $\mathsf{NME\text{-}q\text{-}CCA}_\mathsf{b}^{(1)}$ proceeds like $\mathsf{NME\text{-}q\text{-}CCA}_\mathsf{b}$ except that the DVNIZK proof in the challenge ciphertext is generated by the zero-knowledge simulator for the DVNIZK proof system. To answer the decryption queries, notice that each experiment itself knows all the secret keys, including the DV-NIZK key SP that is required to check the validity of a proof.

If the two experiments are distinguishable, we can construct an adversary that breaks the adaptive zero-knowledge of the DVNIZK. Slightly more precisely, a theorem-chooser/distinguisher pair $(\mathcal{A}_\mathsf{zk}, \mathcal{D}_\mathsf{zk})$ on the DV-NIZK is constructed such that $\mathcal{A}_\mathsf{zk}$ internally simulates the first stage (up to the generation of the challenge ciphertext) of the NME_b experiment, and D_zk internally simulates the second stage. $\mathcal{A}_\mathsf{zk}$ generates all encryption and signature keypairs on its own, but takes the DV-NIZK public key PP from the adaptive zero-knowledge experiment. Since we assume a DV-NIZK with a *strong* adaptive zero-knowledge property, in the corresponding reduction already $\mathcal{A}_\mathsf{zk}$ knows SP and can thus answer decryption queries before the challenge ciphertext is known. This is the only difference from the proof of Claim 1 in [24].

In Claim 2 of [24], the probability for the event $\textsc{BadNIZK}(\mathsf{Expt})$ that the adversary breaks the soundness of the DV-NIZK (in $\mathsf{Expt} \in \{\mathsf{NME}_b, \mathsf{NME}_b^{(1)}, \mathsf{NME}_b^{(2)}\}$) must be shown negligible. For $\mathsf{Expt} = \mathsf{NME}_b$, this is done by constructing an adversary $\mathcal{A}_s$ on the soundness property of the DV-NIZK. Here, $\mathcal{A}_s$ internally simulates the complete NME_b experiment (except for the final decryption of the forged ciphertext vector) and generates all keypairs *except* the DV-NIZK key on its own. The DV-NIZK public key PP is taken from the soundness experiment; since in the [24] CPA setting, no decryptions are necessary, this is sufficient. However, in our q-CCA setting, $\mathcal{A}_s$ might need to answer up to q decryption queries in the $\mathsf{NME\text{-}q\text{-}CCA}$ experiment, and thus needs to check the validity of up to q DV-NIZK proofs. Fortunately, this is exactly what an adversary against the assumed q-adaptive soundness property can do by using the "verifier-oracle" that checks the validity of proofs at most q times.

Then, $\Pr\left[\,\mathsf{NME\text{-}q\text{-}CCA}_\mathsf{b}^{(1)}\,\right] \approx \Pr\left[\,\mathsf{NME\text{-}q\text{-}CCA}_\mathsf{b}\,\right]$, follows similarly (only now by a reduction on the *strong* adaptive zero-knowledge property as before).

The experiment $\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(2)}$ is defined similarly to [24]. However, we cannot show $\Pr\left[\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(1)}\right] = \Pr\left[\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(2)}\right]$, but *can* only show $\Pr\left[\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(1)}\right] \approx \Pr\left[\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(2)}\right]$, which sufficient for the further argument. The reason that we cannot show equality is that the view of an adversary in the $\Pr\left[\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(\mathrm{i})}\right]$ experiments is identical for $i = 1, 2$ only under the condition that the answers to CCA decryption queries do not differ (for $i = 1, 2$; note that in experiment $\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(2)}$, decryption is performed differently than in $\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(1)}$). However, such decryption queries are answered differently only if event BadNIZK happens or if the adversary successfully forged a signature. The probability that one of these events occurs in $\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(1)}$ is negligible, and thus $\Pr\left[\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(1)}\right] \approx \Pr\left[\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(2)}\right]$ follows.

If the adversary succeeds in $\mathsf{NME\text{-}q\text{-}CCA}_{\mathsf{b}}^{(2)}$, we can construct another adversary that breaks the semantic security of the underlying cryptosystem. The rest of the proof is completely analogous to that in [24].

6 Separating NME-CPA from IND-q-CCA

In this section, we show that under bounded chosen ciphertext attacks, non-malleability of the encryption scheme is not immediately implied by indistinguishability. In particular, for any polynomial q, we exhibit an encryption scheme that is IND-q-CCA-secure but is *not non-malleable* under even a chosen plaintext attack (i.e., a malleability attack where the adversary makes no decryption queries). In contrast, it has been shown that unbounded IND-CCA security implies non-malleability (even against unbounded CCA attacks) [10, 1].

$\mathsf{Gen}'(1^k)$: Run $\mathsf{Gen}(1^k)$ and get a pair of keys (pk, sk). Suppose sk is an ℓ-bit string. Choose a random degree-q polynomial $p(x) = p_q x^q + \cdots + p_1 x + sk$ with coefficients in $GF(2^\ell)$ and whose constant term is sk. Output $pk' = pk$ and $sk' = (sk, p)$.
$\mathsf{Enc}'(pk, m)$: Get $c \leftarrow \mathsf{Enc}(pk, m)$ and output $(0, c)$.
$\mathsf{Dec}'(sk, c)$: Parse c as (c_1, c_2). If $c_1 = 0$, output $\mathsf{Dec}(sk, c_2)$. Else, if $c_2 > 0$, output $p(c_2)$ and otherwise return 0.

Fig. 3. An IND-q-CCA-secure PKE scheme PKE′ which is malleable.

Theorem 7. *If there exists an* IND-q-CCA *secure cryptosystem* PKE*, then there exists another* IND-q-CCA *secure cryptosystem* PKE′ *that is not* NME-CPA*-secure.*

REMARK. Theorem 4 shows that the existence of a semantically-secure cryptosystem implies the existence of an IND-q-CCA cryptosystem. Therefore, the "if" clause of the above theorem can be simplified. However, we choose to present it as above to highlight the point that IND-CCA does not imply NME-CPA.

Proof. Assume that there exists an encryption scheme $\text{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ that iss IND-q-CCA-secure. Then, we construct an encryption scheme $\text{PKE}' = (\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}')$ (given in Figure 6) that is also IND-q-CCA-secure, but is not NME-CPA-secure. The proof follows from the two claims shown below.

Claim. $(\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}')$ is IND-q-CCA-secure.

Proof. Suppose that the claim does not hold. We use the adversary $\mathcal{A}$ that breaks the security of $\text{PKE}' = (\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}')$ to construct a q-bounded IND-q-CCA attack against $\text{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. The new adversary $\mathcal{A}'$, on input pk, simply runs $\mathcal{A}(pk)$. When asked to decrypt a ciphertext $(0, c)$, it forwards the query to its own decryption oracle. When asked to decrypt a ciphertext of the form $(1, c_2)$, it returns either 0 if $c_2 = 0$ or a random value. Since $\mathcal{A}$ makes at most q queries, then $\mathcal{A}'$ will be able to answer all queries. The simulation is perfect because the degree-q polynomial $p(\cdot)$ is q-wise independent. This adversary $\mathcal{A}'$ succeeds with the same probability as $\mathcal{A}$, which contradicts the assumption that PKE is q-bounded secure. □

Claim. $(\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}')$ is not NME-CPA-secure.

Proof. Without loss of generality, assume that the message space of PKE include the bits 0 and 1. On input a public key pk, the adversary submits as a message pair, 0 and 1.

Upon receiving a ciphertext c, the attacker first computes $\alpha = \mathsf{Enc}(pk, c)$. It then returns the vector $(\alpha, \beta_1, \ldots, \beta_{q+1})$ where $\beta_i = (1, i)$.

Notice that the output of the experiment is the vector $(c, p(1), \ldots, p(q+1))$. The distinguisher D now works as follows. It first uses $p(1), \ldots, p(q+1)$ to interpolate the secret key sk, and then runs $\mathsf{Dec}(sk, c)$ and prints the result as its output.

The distinguisher's output in the NME_0 experiment will therefore be 0 and its output in the NME_1 will be 1, which shows that PKE' is not even NME-CPA-secure.

As one final point, it may be that the message space of PKE does not include the ciphertext — for example, the size of the ciphertext may be too big. This is easily handled. The adversary can simply encode c in a bit-by-bit fashion over many ciphertexts, and the distinguisher can simply reconstruct c to perform its test. □

Acknowledgments

We thank Ivan Damgård, Tal Malkin, and Moti Yung for their comments.

References

1. Mihir Bellare, Anand Desai, David Pointcheval, and Phillip Rogaway. Relations among notions of security for public-key encryption schemes. In CRYPTO '98, volume 1462 of LNCS, pages 26–45. Springer-Verlag, Berlin, Germany, August 1998.
2. Ran Canetti, Yevgeniy Dodis, Shai Halevi, Eyal Kushilevitz, and Amit Sahai. Exposure-resilient functions and all-or-nothing transforms. In EUROCRYPT 2000, volume 1807 of LNCS, pages 453–469. Springer-Verlag, Berlin, Germany, May 2000.
3. Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In EUROCRYPT 2004, volume 3027 of LNCS, pages 207–222. Springer-Verlag, Berlin, Germany, May 2004.
4. Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In CRYPTO' 94, volume 839 of LNCS, pages 174–187. Springer-Verlag, Berlin, Germany, August 1994.
5. Ronald Cramer, Dennis Hofheinz, and Eike Kiltz. A note on bounded chosen ciphertext security from black-box semantical security. Cryptology ePrint Archive, Report 2006/391, 2006. `http://eprint.iacr.org/`.
6. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In CRYPTO '98, volume 1462 of LNCS, pages 13–25. Springer-Verlag, Berlin, Germany, August 1998.
7. Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In EUROCRYPT 2002, volume 2332 of LNCS, pages 45–64. Springer-Verlag, Berlin, Germany, April / May 2002.
8. Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM Journal on Computing*, 33(1):167–226, 2003.
9. Yevgeniy Dodis, Jonathan Katz, Shouhuai Xu, and Moti Yung. Key-insulated public key cryptosystems. In EUROCRYPT 2002, volume 2332 of LNCS, pages 65–82. Springer-Verlag, Berlin, Germany, April / May 2002.
10. Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.
11. P. Erdös, P. Frankel, and Z. Furedi. Families of finite sets in which no set is covered by the union of r others. *Israeli Journal of Mathematics*, 51:79–89, 1985.
12. Yael Gertner, Tal Malkin, and Steven Myers. Towards a separation of semantic and CCA security for public key encryption. In TCC 2007, LNCS 4392, pages 434–455. Springer-Verlag, Berlin, Germany, 2007.
13. Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004.
14. Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984.
15. Goichiro Hanaoka and Hideki Imai. A generic construction of CCA-secure cryptosystems without NIZKP for a bounded number of decryption queries. Cryptology ePrint Archive, Report 2006/408, 2006. `http://eprint.iacr.org/`.
16. Swee-Huay Heng and Kaoru Kurosawa. k-resilient identity-based encryption in the standard model. In CT-RSA 2004, volume 2964 of LNCS, pages 67–80. Springer-Verlag, Berlin, Germany, February 2004.
17. Ravi Kumar, Sridhar Rajagopalan, and Amit Sahai. Coding constructions for blacklisting problems without computational assumptions. In CRYPTO'99, volume 1666 of LNCS, pages 609–623. Springer-Verlag, Berlin, Germany, August 1999.

18. Kaoru Kurosawa and Yvo Desmedt. A new paradigm of hybrid encryption scheme. In CRYPTO 2004, volume 3152 of LNCS, pages 426–442. Springer-Verlag, Berlin, Germany, August 2004.
19. L. Lamport. Constructing digital signatures from a one-way function. Technical Report CSL-98, SRI International, october 1979.
20. Yehuda Lindell. A simpler construction of CCA2-secure public-key encryption under general assumptions. *Journal of Cryptology*, 19(3):359–377, July 2006.
21. Moni Naor and Moti Yung. Universal one-way hash functions and their cryptographic applications. In STOC 89, pages 33–43. ACM Press, May 1989.
22. Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In STOC 90. ACM Press, May 1990.
23. Rafael Pass, Abhi Shelat, and Vinod Vaikuntanathan. Bounded-CCA secure non-malleable encryption. MIT CSAIL Technical Report TR-2006-081. December 2006.
24. Rafael Pass, Abhi Shelat, and Vinod Vaikuntanathan. Construction of a non-malleable encryption scheme from any semantically secure one. In CRYPTO 2006, volume 4117 of LNCS, pages 271–289. Springer-Verlag, Berlin, Germany, August 2006.
25. Duong Hieu Phan and David Pointcheval. About the security of ciphers (semantic security and pseudo-random permutations). In SAC 2004, volume 3357 of LNCS, pages 182–197. Springer-Verlag, Berlin, Germany, August 2004.
26. Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In CRYPTO'91, volume 576 of LNCS, pages 433–444. Springer-Verlag, Berlin, Germany, August 1992.
27. Rompel, J. One-way Functions are necessary and sufficient for secure signatures. In STOC '90, pages 387 – 394.
28. Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In FOCS '99, pages 543–553. IEEE Computer Society Press, October 1999.