A Note on Bounded Chosen Ciphertext Security from Black-box Semantical Security

Ronald Cramer[1,2] Dennis Hofheinz[1] Eike Kiltz[1]

[1] CWI Amsterdam
The Netherlands
{cramer,hofheinz,kiltz}@cwi.nl

[2] Mathematical Institute
Leiden University
The Netherlands

Abstract

Designing public key encryption schemes withstanding chosen ciphertext attacks, which is the highest security level for such schemes, is generally perceived as a delicate and intricate task, and for good reason. In the standard model, there are essentially three well-known but quite involved approaches. This state of affairs is to be contrasted with the situation for semantically secure encryption schemes, a much weaker security notion that only guarantees security in the absence of active attack but that appears to be much easier to fulfill, both conceptually and practically. Thus, the boundary between passive attack and active attack seems to make up the dividing line between which security levels are relatively easily achieved and which are not. Our contributions are two-fold.

First, we show a simple, efficient black-box construction of a public key encryption scheme withstanding chosen ciphertext attack from any given semantically secure one. Our scheme is q-bounded in the sense that security is only guaranteed if the adversary makes at most q adaptive chosen ciphertext queries. Here, q is an arbitrary polynomial that is fixed in advance in the key-generation. Our work thus shows that whether or not the number of active, adversarial queries is known in advance is the dividing line, and not passive versus active attack. In recent work, Gertner, Malkin and Myers show that such black-box reductions are impossible if instead q is a polynomial that only depends on the adversary. Thus, in a sense, our result appears to be the best black-box result one can hope for. Second, we give a non-blackbox reduction from bounded chosen ciphertext security to semantic security where the length of the public/secret keys and ciphertexts drops from quadratic to linear in q, compared to our black-box construction. This latter scheme, however, is only of theoretical interest as it uses general NP-reductions, and our blackbox construction is in fact much more practical.

Keywords: Black-box construction, chosen-ciphertext security

1 Introduction

Designing public key encryption schemes withstanding chosen ciphertext attacks, which is the highest security level for such schemes, is generally perceived as a delicate and intricate task, and for good reason. In the standard model, there are essentially three approaches known. The first approach, pioneered by Naor and Yung [15] in the early 1990s, and subsequently extended by Dolev, Dwork and Naor [7], and later Sahai [20] and Lindell [14], is based on the use of non-interactive zero knowledge for NP. This leads to schemes based on quite general

cryptographic assumptions that are of theoretical significance only. The second is due to Cramer and Shoup [3, 4, 5] and is based on hash-proof systems. This leads to quite practical schemes based on several concrete number-theoretical assumptions. The third and most recent approach is due to Canetti, Halevi and Katz [2], and relies on identity-based cryptography.

This state of affairs is to be contrasted with the situation for semantically secure encryption schemes, a much weaker security notion that only guarantees security in the absence of active attack: the design of such schemes appears to be a much easier task, given the relative abundance of quite practical schemes that have been shown to exists under several different assumptions, including general assumptions such as trapdoor one-wayness [10]. Thus, the boundary between passive attack and active attack seems to make up the dividing line between which security levels are relatively easily achieved and which are not. Our contributions are two-fold.

First, we show a simple, efficient black-box construction of a public key encryption scheme withstanding chosen cipher-text attack from any given semantically secure one. Our scheme is q-bounded in the sense that security is only guaranteed if the adversary makes at most q adaptive chosen cipher-text queries. Here, q is an arbitrary polynomial that is fixed in advance in the key-generation. Technically, our result is a combination of techniques from [2, 6]. 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.

Our work thus shows that whether or not the number of active, adversarial queries is known in advance is the dividing line, and not passive versus active attack. In recent work, Gertner, Malkin and Myers [9] show that such black-box reductions are impossible if instead q is a polynomial that only depends on the adversary. Thus, in a sense, our result appears to be the best black-box result one can hope for.

Second, the size of the public key, that of the secret key and the size of the cipher-text all depend quadratically on the bound q. Building on recent work by Pass, Shelat and Vaikuntanathan [17], we also give a non-blackbox reduction from bounded chosen cipher-text security to semantical security where this dependence is linear. This latter scheme, however, is only of theoretical interest as it uses general NP-reductions.

1.1 Black-Box reductions

One natural task in modern cryptography is to relate different primitives to eachother. In such reductions we assume that some primitive $\mathcal{P}$ exists and we want to infer that some different primitive $\mathcal{Q}$ also exists, i.e. that primitive $\mathcal{Q}$ can be constructed from $\mathcal{P}$. If in the construction of $\mathcal{Q}$ there are only oracle-calls to the primitive $\mathcal{P}$ and no special structure of $\mathcal{P}$ is used, we speak of a black-box construction. If $\mathcal{Q}$ can be construction from $\mathcal{P}$ in a black-box way, we write $\mathcal{Q}^{\mathcal{P}}$. Almost all known constructions in cryptography are indeed black-box (such as the equivalence of one-way functions and digital signatures [19, 13, 10]).

In terms of negative results, Impagliazzo and Rudich initiated a line of research showing that certain block-box reductions cannot exist. In particular [12] shows a black-box separation between key agreement and one-way functions. Recently, Gertner, Malkin and Myers [9] show a certain black-box separation between chosen-ciphertext secure encryption and semantically secure encryption schemes. More formally, they prove the following.

Theorem [9] There exists no black box reduction that from a given semantically secure $\mathcal{PKE} = (\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ constructs a chosen-ciphertext secure $\mathcal{PKE} = (\mathsf{KG}^{\mathsf{kg},\mathsf{enc},\mathsf{dec}}, \mathsf{ENC}^{\mathsf{kg},\mathsf{enc},\mathsf{dec}}, \mathsf{DEC}^{\mathsf{kg},\mathsf{dec}})$.

This is contrasted by our main result.

Main Theorem For any fixed polynomial $q(k)$, there exists a black-box reduction that from a given semantically secure $\mathcal{pke} = (\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ constructs a $q(k)$ bounded chosen-ciphertext secure $\mathcal{PKE} = (\mathsf{KG}^{\mathsf{kg}}, \mathsf{ENC}^{\mathsf{kg},\mathsf{enc}}, \mathsf{DEC}^{\mathsf{kg},\mathsf{dec}})$.

Note that the black-box separation result from [9] only holds for reductions where decryption DEC itself does not make calls to enc (re-encryption). Our black-box reduction of q bounded CCA encryption falls exactly into this category, i.e. our construction does not use re-encryption.

1.2 Non black-box constructions

While most of the known reductions in cryptography are black-box, there are some interesting reductions that ore non black-box, i.e. the construction of $\mathcal{Q}$ from $\mathcal{P}$ may make use of a particular structure of $\mathcal{P}$ (for example, $\mathcal{P}$'s circuit representation). Most importantly, all constructions of chosen-ciphertext secure encryption from generic assumptions (such as the existence of enhanced trapdoor permutations) are non black-box [7, 20, 14]. This also includes the more recent result by Pass, Shelat and Vaikuntanathan [17] who give a non-blackbox reduction from non-malleability to semantical security, without any further complexity theoretic assumption. While the size of the public key from our black-box q bounded chosen-ciphertext construction was quadratic in q, we build on [17] to improve this result using non back-box techniques.

Theorem For any fixed polynomial $q(k)$, there exists a (non-black-box) reduction that from any given semantically secure PKE scheme constructs a $q(k)$-bounded chosen ciphertext secure PKE scheme. The size of the public key in that construction is linear in $q(k)$.

The above non black-box construction even reaches the stronger security level of $q(k)$-bounded chosen-ciphertext non-malleability. We remark that, even though the parameters of our non black-box construction only depend linearly on $q(k)$, due to the use of generic NP-reductions its overall complexity only compares favorably to our black-box construction for very large polynomials $q(k)$.

1.3 Organization

After fixing some notation in Section 2 we formally define $q(k)$-bounded chosen-ciphertext security for PKE scheme in Section 3. In Section 4 we formally state our main result and provide a proof by presenting our black-box constructions. Finally, Section 5 deals with the mentioned non black-box extensions.

1.4 A remark on related work

We have recently sent a copy of our note to the authors of [17]. In immediate return, we received an unpublished manuscript [16] of theirs whose main result and techniques are essentially identical to our subresult from Section 5, i.e., a non black-box reduction from q-bounded chosen-ciphertext security to semantic security. They have also pointed out to us that at the very end of their presentation of [17] at CRYPTO'06 they have announced this non black-box result, including a very brief indication on a guiding observation.

We have achieved our subresult on non black-box reductions from Section 5 entirely independently, and only motivated by a (theoretical) efficiency issue prompted by our main result from Section 4, i.e., our *black-box* reduction from q-bounded chosen-ciphertext security to semantic security. Nevertheless, we do believe that the main result from their manuscript [16] predates our non black-box result from Section 5. We stress, however, that their manuscript [16] does not claim anything similar to our main result.

2 Notation

If x is a string, then $|x|$ denotes its length, while if S is a set then $|S|$ denotes its size. If $k \in \mathbb{N}$ then 1^k denotes the string of k ones. If S is a set then $s \xleftarrow{\$} S$ 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 $\mathcal{O}_1, \mathcal{O}_2, \ldots$ and by $z \xleftarrow{\$} \mathcal{A}^{\mathcal{O}_1, \mathcal{O}_2, \cdots}(x, y, \ldots)$ we denote the operation of running $\mathcal{A}$ with inputs $(x, y, \ldots)$ and Black-box access to oracles $\mathcal{O}_1, \mathcal{O}_2, \ldots$, and letting z be the output.

3 Public-Key Encryption

Definition 3.1 A triple $\mathcal{pke} = (\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ is a public-key encryption (PKE) scheme, if kg and enc are probabilistic PTA, and dec is a deterministic polynomial-time algorithm. For consistency, we require that for all $k \in \mathbb{N}$, all messages M, it must hold that $\Pr[\mathsf{dec}(sk, \mathsf{enc}(pk, M)) = M] = 1$, where the probability is taken over the above randomized algorithms and $(pk, sk) \xleftarrow{\$} \mathsf{kg}(1^k)$.

Definition 3.2 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

$$\mathbf{Adv}^{ind\text{-q-cca}}_{\mathcal{PKE},\mathcal{A}}(k) \;=\; \Pr[\mathbf{Exp}^{ind\text{-q-cca-1}}_{\mathcal{PKE},\mathcal{A}}(k) = 1] - \Pr[\mathbf{Exp}^{ind\text{-q-cca-0}}_{\mathcal{PKE},\mathcal{A}}(k) = 1]$$

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

$$\begin{array}{l}
\textbf{Experiment } \mathbf{Exp}^{ind\text{-q-cca-b}}_{\mathcal{PKE},\mathcal{A}}(k) \\
\quad (pk, sk) \xleftarrow{\$} \mathsf{KG}(1^k) \\
\quad (M_0, M_1, St_1) \xleftarrow{\$} \mathcal{A}_1^{\mathsf{DEC}(sk,\cdot)}(pk) \text{ s.t. } |M_0| = |M_1| \\
\quad c^* \xleftarrow{\$} \mathsf{ENC}(pk, M_b) \\
\quad b' \xleftarrow{\$} \mathcal{A}_2^{\mathsf{DEC}(sk,\cdot)}(c^*, St_1) \\
\quad \text{If } b \neq b' \text{ then return 0 else return 1.}
\end{array}$$

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 in the second stage may contain c^*. PKE scheme $\mathcal{PKE}$ is said to be indistinguishable against q bounded chosen-ciphertext attacks (IND-q-CCA secure in short) if the advantage function $\mathbf{Adv}^{ind\text{-q-cca}}_{\mathcal{PKE},\mathcal{A}}(k)$ is a negligible function in k for all adversaries $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ with PPT $\mathcal{A}_1$, $\mathcal{A}_2$.

We have the following relation to the standard security definitions for PKE schemes. Scheme $\mathcal{PKE}$ is said to be

- indistinguishable against chosen-plaintext attacks [11] (semantically secure or IND-CPA) if it is IND-0-CCA secure.
- indistinguishable against chosen-ciphertext attacks [18] (IND-CCA) if it is IND-q-CCA secure for any polynomial $q(k)$.

4 Black-box construction of bounded CCA secure encryption

4.1 Overview

The general outline of our construction is as follows. First, as demonstrated by Canetti, Halevi, and Katz [2], every identity-based encryption scheme can be transformed into a chosen-ciphertext secure PKE scheme. Second, a semantically secure PKE scheme implies a "$q(k)$-resilient" identity-based encryption scheme. (The notion of q-resilient security in the context of identity-based encryption 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 [6]. A closer observation of the combination of the two results already reveals the construction of our $q(k)$-bounded chosen-ciphertext secure PKE scheme. Since both transformations are black-box our main result can be obtained. 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 4.1 For any fixed $q(k)$, there exists a black-box reduction that from a given semantically secure $(\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ constructs a q bounded chosen-ciphertext secure PKE $(\mathsf{KG}^{\mathsf{kg}}, \mathsf{ENC}^{\mathsf{kg},\mathsf{enc}}, \mathsf{DEC}^{\mathsf{kg},\mathsf{dec}})$.

Here we give direct proof of this theorem that bypasses the notion of identity-based encryption alltogether.

4.2 Building blocks

Cover-free families. We start by recalling cover-free families. If S, T are sets, we say that S does not cover T if $T \not\subseteq S$. 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 [8]: there is a deterministic PTA 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 PTA that on input s, q, t returns F_t. We call $F_t = \mathsf{SUB}(s(k), q(k), t)$ the subset associated to index $t \in \{1, \ldots, s(k)\}$.

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. We need a one-time signature scheme $\mathcal{OTS} = (\mathsf{sigkg}, \mathsf{sign}, \mathsf{vfy})$ [13]. We require that this scheme be secure in the sense of strong unforgeability against one-time attacks. We assume that the verification keys which are part of the output by sigkg are bitstrings of size k which we interpret as natural numbers in $\{1, \ldots, 2^k\}$. One-time signatures can be constructed in a black-box way from semantically secure encryption. This follows by combining the observation that semantically secure encryption implies one-way functions, the fact that one-way functions imply universal one-way hash functions [19], and the result that universal one-way hash functions imply strong one-time signatures [13, 10]. All transformations are black-box. The implied $\mathcal{OTS}$

Algorithm $\mathsf{KG}^{\mathsf{kg}}(1^k)$
Define $s(k) = 2^k, d(k) = 16kq^2(k), l(k) = 4kq(k)$ as in Equation (1)
for $i = 1, \ldots, d(k)$ do $(pk_i, sk_i) \xleftarrow{\$} \mathsf{kg}(1^k)$.
$PK \leftarrow (pk_1, \ldots, pk_{d(k)})$; $SK \leftarrow (sk_1, \ldots, sk_{d(k)})$
Return (PK, SK)

Algorithm $\mathsf{ENC}^{\mathsf{kg},\mathsf{enc}}(PK, M)$
$(verk, sigk) \xleftarrow{\$} \mathsf{sigkg}^{\mathsf{kg}}(1^k)$
Let $F_{verk} = \{s_1, \ldots, s_{l(k)}\}$ be the subset associated to $verk$
Pick random $M_1, \ldots, M_{l(k)}$ s.t. $M = \bigoplus_{i=1}^{l(k)} M_i$
For $j = 1, \ldots, l(k)$ do $c_j \xleftarrow{\$} \mathsf{enc}(pk_{s_j}, M_j)$
$c \leftarrow (c_1, \ldots, c_{l(k)})$
$\sigma \xleftarrow{\$} \mathsf{sign}^{\mathsf{kg}}(sigk, c)$
Return $C \leftarrow (c, verk, \sigma)$

Algorithm $\mathsf{DEC}^{\mathsf{kg},\mathsf{dec}}(SK, C)$
Parse C as $(c, verk, \sigma)$
If $\mathsf{vfy}^{\mathsf{kg}}(verk, c, \sigma) = \perp$ then return $\perp$.
Let $F_{verk} = \{s_1, \ldots, s_{l(k)}\}$ be the subset associated to $verk$
Parse c as $(c_1, \ldots, c_{l(k)})$
For $j = 1, \ldots, l(k)$ do $M_j \leftarrow \mathsf{dec}(sk_{s_j}, c_j)$
$M \leftarrow M_1 \oplus \ldots \oplus M_{l(k)}$
Return M

Figure 1: Black-box construction of an IND-q-CCA secure PKE scheme $\mathcal{PKE} = (\mathsf{KG}^{\mathsf{kg}}, \mathsf{ENC}^{\mathsf{kg},\mathsf{enc}}, \mathsf{DEC}^{\mathsf{kg},\mathsf{dec}})$ from any semantically secure PKE scheme $\mathcal{pke} = (\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$

algorithms make black-box calls to evaluations of the one-way function which itself makes black-box calls to kg.[1] Hence, there exists a black-box reduction that from a given semantically secure $(\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ constructs a strongly unforgeable one-time signature $\mathcal{OTS} = (\mathsf{sigkg}^{\mathsf{kg}}, \mathsf{sign}^{\mathsf{kg}}, \mathsf{vfy}^{\mathsf{kg}})$.

4.3 The construction

Let $q(k) : \mathbb{N} \to \mathbb{N}$ be a function. We build a PKE scheme $\mathcal{PKE} = (\mathsf{KG}, \mathsf{ENC}, \mathsf{DEC})$ with black-box access to $\mathcal{pke} = (\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ as follows.

$\mathcal{PKE}$ key generation generates d independent instances of public/secret key pairs of $\mathcal{pke}$. Encryption first creates a random key-pair of the one-time signature scheme to create a verification/signing key. The resulting verification key uniquely refers to a subset $F_{verk} = \{s_{i_1}, \ldots s_{i_l}\} \subseteq \{1, \ldots, d(k)\}$ of the base set of the cover-free family. The message M is shared using an all-or-nothing transform into random l parts subject to $M = M_1 \oplus \ldots \oplus M_l$, and each share M_i is encrypted using public key pk_{s_i}. The final ciphertext consists of the one-time verification key, the l concatenated ciphertexts, and a signature of the latter concatenated ciphertexts using the one-time signing key. We stress that encryption is stateless. Decryption first checks the signature using the verification key contained in the ciphertext. If the signature is valid, it reverses the encryption process using the set of secret keys $(sk_i)_{1 \le i \le l}$, where $F_{verk} = \{s_{i_1}, \ldots s_{i_l}\}$ is again uniquely derived from the verification key. Finally, combining the resulting plaintexts by XOR yields the message.

A more formal description is given in Figure 1. In general we can also use any computationally secure AONT (e.g., the black-box construction from [1] based on oneway functions) to decrease ciphertext size.

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

The following proves our main result, Theorem 4.1.

[1] The oneway function $f_k(\cdot)$ is defined as $f_k(r) = pk$, where $(pk, sk) \leftarrow \mathsf{kg}(1^k, r)$.

Claim 4.2 If $\mathcal{pke} = (\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ is semantically secure then $\mathcal{PKE} = (\mathsf{KG}^{\mathsf{kg}}, \mathsf{ENC}^{\mathsf{kg},\mathsf{enc}}, \mathsf{DEC}^{\mathsf{kg},\mathsf{dec}})$ as described in Figure 1 is indistinguishable against q bounded chosen-ciphertext attacks.

Proof: (Sketch.) Assume there exists an adversary $\mathcal{A}$ against the IND-q-CCA security of $\mathcal{PKE}$. We show that then there exists either an adversary $\mathcal{B}$ against the semantical security security of $\mathcal{pke}$ or an adversary $\mathcal{C}$ against the strong unforgeability of $\mathcal{OTS}$.

We first describe adversary $\mathcal{B}$.

Setup. Given security parameter k, adversary $\mathcal{B}$ sets up a random instance $(verk^*, sigk^*)$ of the signature scheme via $\mathsf{sigkg}(1^k)$. Let F_{verk^*} be the subset associated to index $verk^*$, where $F_{verk^*} = \{f_1^*, \ldots, f_{l(k)}^*\}$. Adversary $\mathcal{B}$ picks a random

$$f^* \xleftarrow{\$} F_{verk^*} \ . \tag{2}$$

$\mathcal{B}$ generates independent pairs of keys (pk_i, sk_i) for $i \in \{1, \ldots, d(k)\} \setminus \{f^*\}$. The public key PK is as $(pk_1, \ldots, pk_d)$, where the "rogue key" pk_{f^*} is the public key provided by $\mathcal{B}$'s experiment for semantical security of $\mathcal{pke}$. Hence, $\mathcal{B}$ knows all secret keys $sk_1, \ldots, sk_{d(k)}$, except sk_{f^*}.

Adversary $\mathcal{A}$ is run on PK answering to its queries as follows.

Challenge query. Adversary $\mathcal{B}$ will generate $\mathcal{A}$'s challenge ciphertext with respect to the signature pairs $(verk^*, sigk^*)$ generated during setup. Adversary $\mathcal{B}$ receives the two messages M_0^*, M_1^* from $\mathcal{A}$ and, for $1 \le i \le l(k)$, picks random messages M_i. Let $M = \bigoplus_{i=1}^{l(k)} M_i$ and let $i^* \in \{1, \ldots, l(k)\}$ such that $f^* = f_{i^*}$. For $c \in \{0,1\}$ define $M_{i^*,c} = M_{i^*} \oplus M \oplus M_c^*$ such that $M_c^* = M_{i^*,c} \oplus \bigoplus_{i \neq i^*} M_i$. Adversary $\mathcal{B}$ forwards $M_{i^*,0}, M_{i^*,1}$ to the surrounding semantical security experiment and receives a $\mathcal{pke}$ challenge ciphertext $c_{i^*}^*$ for message $M_{i^*,b}$ for unknown b. Then $\mathcal{A}$'s ciphertext vector $c^* = (c_1^*, \ldots, c_{l(k)}^*)$ is filled for $i \neq i^*$ by encrypting $c_i^* \xleftarrow{\$} \mathsf{enc}(pk_{f_i^*}, M_i)$. Finally, $c^* = (c_1^*, \ldots, c_{l(k)}^*)$ is signed using the one-time signing key $sigk^*$ to obtain σ^*. Adversary $\mathcal{A}$'s challenge ciphertext is $C^* = (c^*, verk^*, \sigma^*)$. Note that C^* is a correctly distributed $\mathcal{PKE}$ ciphertext for $\mathcal{B}$'s challenge message M_b^*.

Decryption queries. Suppose $\mathcal{A}$ makes a decryption query $C = (c, verk, \sigma)$ containing a valid signature σ (satisfying $\mathsf{vfy}(verk, c) \neq \perp$). Let F_{verk} be the subset associated to index $verk$. We distinguish between the following mutually exclusive cases:

1. $verk = verk^*$. If $(c, \sigma) = (c^*, \sigma^*)$ then $\mathcal{A}$ made an illegal query with $C = C^*$. If $(c, \sigma) \neq (c^*, \sigma^*)$ then $\mathcal{A}$ has broken the strong unforgeability of the one-time signature scheme.
2. $verk \neq verk^*$, $f^* \notin F_{verk}$. In that case $\mathcal{B}$ knows all secret keys sk_i for $i \in F_{verk}$ and hence can correctly decrypt.
3. $verk \neq verk^*$, $f^* \in F_{verk}$. Adversary $\mathcal{B}$ returns a random bit and terminates. Denote this event by FAIL.

Output. Eventually, $\mathcal{A}$ outputs a guess bit b'. $\mathcal{B}$ outputs the same bit b' and terminates.

This completes the description of $\mathcal{B}$.

If FAIL does not happen $\mathcal{B}$ has the same probability of winning the semantical security experiment as $\mathcal{A}$. We claim that $\Pr[\neg\text{FAIL}]$ is bounded by $1/l(k)$. For $1 \le i \le q$, let $verk_i$ be

the verification key from $\mathcal{A}$'s ith decryption query. By the properties of the q cover-free family we have that $F_{verk^*} \not\subseteq \bigcup_{i=1}^{q} F_{verk_i}$, that means that there exists an element $f \in F_{verk^*}$ such that $f \notin \bigcup_{i=1}^{q} F_{verk_i}$. As long as $\mathcal{B}$ picked this $f^* = f$ in Equation (2) at the beginning of the experiment from the set F_{verk^*} of cardinality $l(k)$, event FAIL does not happen. Now the claim follows since $\mathcal{A}$'s simulation is independent of the choice of f^*. ▌

Remark 4.3 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 $\mathcal{PKE}$ from semantically secure $\mathcal{pke}$ 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_{verk}, where $verk \in \{1, \ldots, q(k)\}$ is one of the $q(k)$ distinct public keys of $\mathcal{pke}$, and $verk$ 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 as follows. First guess $(verk^*, sigk)$, where $verk^*$ is the verification key used for the (yet unknown) challenge ciphertext. Then make decryption queries of the form $(c, verk^*, \mathsf{sign}(sigk, c))$ for arbitrary $\mathcal{pke}$ ciphertexts c. This way, with probability $1/q(k)$, any CCA attack to the original scheme $\mathcal{pke}$ translates to a CCA attack of $\mathcal{PKE}$; if $\mathcal{pke}$ gets insecure after one single CCA1 decryption query so does $\mathcal{PKE}$.

Remark 4.4 It might be interesting to explore what (additional) security properties $\mathcal{PKE}$ satisfies once invoked with a scheme $\mathcal{pke}$ that itself is not only IND-CPA secure, but, say, NM-CPA secure. Unfortunately, we cannot hope that $\mathcal{PKE}$ is NM-CPA secure, independent of $\mathcal{pke}$'s security: say that adversary $\mathcal{A}$ receives a challenge ciphertext $C^* = (c^*, verk^*, \sigma^*)$ with $c^* = (c_1, \ldots, c_l)$ and $F_{verk^*} = \{s_1^*, \ldots, s_l^*\}$. 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 $s_i^* \in F^{(i)} \neq F_{\text{verk}}$, 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.

On the other hand, IND-CCA1 security of $\mathcal{pke}$ implies that $\mathcal{PKE}$ is secure against IND-attackers that have full access to a decryption oracle in the first phase of the attack (i.e., before receiving the challenge ciphertext) but only limited access (limited to q queries) to it in the second attack phase. (The reduction is essentially the same as the one above.)

5 A non-black-box bounded CCA secure construction

5.1 Overview

If one is willing to sacrifice the black box property of the construction from the previous section, a significant gain in efficiency is possible. Namely, we work with the construction from [17]. This construction takes any IND-CPA secure public key encryption scheme $\mathcal{pke}$ and transforms it (in a non-black-box way) into a scheme $\mathcal{nmpke}$ that is non-malleable under chosen-plaintext attacks. This transformation first constructs a certain type of designated-verifier non-interactive zero knowledge (DV-NIZK) proof system from $\mathcal{pke}$ and then employs the paradigm from [15, 7] to achieve non-malleability.

However, as [17] points out, $\mathcal{nmpke}$ may generally not be CCA secure. The reason is that the constructed DV-NIZK proof system used in $\mathcal{nmpke}$ is insecure under (sequential) composition. However, we show that already a slight tweak in their DV-NIZK construction suffices to achieve precisely the form of composability needed to prove $\mathcal{nmpke}$ IND-q-CCA secure. This results in an overall construction of an IND-q-CCA scheme with a public key of size *linear* in q.

Formally, we show the following:

Theorem 5.1 For any fixed polynomial $q(k)$, there exists a (non-black-box) reduction that from any given semantically secure $(\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ constructs an IND-q-CCA secure PKE $\mathcal{PKE}$. The size of the public key in that construction $\mathcal{PKE}$ is linear in q.

5.2 The DV-NIZK scheme of [17]

The construction of [17] takes an IND-CPA secure public key encryption scheme $\mathcal{pke} = (\mathsf{kg}, \mathsf{enc}, \mathsf{dec})$ and uses it as follows to generate a DV-NIZK proof system $(\mathcal{D}, P, V)$, where $\ell := k$ and (P_1, V_1, P_2, V_2) is a suitable Σ-protocol for zero-knowledge proofs:

- *Sampling algorithm:* $\mathcal{D}(1^k)$ chooses $f = (f_1, \ldots, f_\ell) \in \{0,1\}^\ell$ and generates 2ℓ $\mathcal{pke}$ keypairs $(pk_{i,j}, sk_{i,j})$ (where $i \in \{1, \ldots, \ell\}$ and $j \in \{0,1\}$). Output is (PP, SP) where

  $$\text{PP} = (pk_{0,1}, pk_{1,1}, \ldots, pk_{0,\ell}, pk_{1,\ell})$$

  and

  $$\text{SP} = (f, sk_{f_1,1}, \ldots, sk_{f_\ell,\ell}).$$

- *Prover:* $P(\text{PP}, x, w)$ generates triples

  $$(a_i, s_i) \leftarrow P_1(x, w)$$
  $$c_{b,i} \leftarrow P_2(s, b)$$
  $$\alpha_{b,i} \leftarrow \mathsf{enc}_{pk_{b,i}}(c_{b,i})$$

  for all $i \in \{1, \ldots, \ell\}$ and $b \in \{0,1\}$. Output of the prover is $\pi := ((a_i, \alpha_{0,i}, \alpha_{1,i}))_{i=1}^{\ell}$.

- *Verifier:* $V(\text{PP}, \text{SP}, x, \pi)$ parses π as $\pi = ((a_i, \alpha_{0,i}, \alpha_{1,i}))_{i=1}^{\ell}$ and checks if for all $i \in \{1, \ldots, \ell\}$ and $m_i := \mathsf{dec}_{sk_{f_i,i}}(\alpha_{f_i,i})$, it holds that $V_2(a_i, f_i, m_i)$ accepts. If all V_2 instances accept, V also accepts, otherwise, V rejects.

In other words, non-interactivity is achieved by "predistributing" the choices for the challenge bits b_i in form of the secret keys $sk_{f_i,i}$. The verifier only knows the secret keys $sk_{f_i,i}$ corresponding to his (fixed in advance) "virtual choices" of $b_i = f_i$. Because these choices are fixed once and for all and cannot be changed during the protocol, the randomness that the verifier may employ is limited in some sense by ℓ. For that reason, there is a simple attack on the soundness property of the protocol, once an adversary may ask questions about the validity of, say, ℓ proofs of his choice. This adversary may take correct proofs for arbitrary valid statements, then substitute one encryption $\alpha_{b,i}$ and will then learn from the (in-)validity of the whole proof whether $f_i = b$ or not.

5.3 Our modification

Note that [17] only consider $\ell = k$, and so the scheme completely breaks down after k such questions. On the other hand, it seems intuitive that if we take $\ell = q + k$, then at least q such questions can be tolerated without giving up the soundness property. It turns out that this is true, and actually all that is needed to make the whole construction of [15, 7, 17] IND-q-CCA secure.

Definition 5.2 (q-adaptive security) Let $q = q(k)$ be a polynomial, and let $(\mathcal{D}, P, V)$ be a designated verifier non-interactive zero-knowledge proof system for an $\mathcal{NP}$-language L with witness relation R_L in the sense of [17]. We say that $(\mathcal{D}, P, V)$ has *q-adaptive security* iff there is a negligible function μ, such that

1. **(q-adaptive soundness.)** For every prover algorithm B and every $k \in \mathbb{N}$,

$$\Pr\left[(\text{PP}, \text{SP}) \leftarrow \mathcal{D}(1^k)\,;\, (x, \pi) \leftarrow B^{\mathcal{O}^q_{\text{PP},\text{SP}}}(\text{PP}) \,:\, x \notin L \text{ but } V(\text{PP}, \text{SP}, x, \pi) = 1 \right] \leq \mu(|x|),$$

   where the oracle $\mathcal{O}^q_{\text{PP},\text{SP}}$ returns $V(\text{PP}, \text{SP}, x, \pi)$ on input (x, w), but only for the first q queries. That is, $\mathcal{O}^q_{\text{PP},\text{SP}}$ provides B with the possibility to check the validity of q adaptively chosen proofs.

2. **(Strong adaptive zero-knowledge.)** For every PPT theorem chooser A, there exists a simulator $S = (S_1, S_2)$ such that the outputs of the following experiments are indistinguishable.

Experiment $\mathsf{ZK}_A(k)$	**Experiment** $\mathsf{ZK}_A(k)$
$(\text{PP}, \text{SP}) \leftarrow \mathcal{D}(1^k)$	$(\text{PP}', \text{SP}', \text{STATE}) \leftarrow S_1(1^k)$
$(x, w, \text{STATE}_A) \leftarrow A(\text{PP}, \text{SP})$	$(x, w, \text{STATE}_A) \leftarrow A(\text{PP}', \text{SP}')$
$\pi \leftarrow P(\text{PP}, x, w)$	$\pi' \leftarrow S_2(\text{PP}', \text{SP}', \text{STATE})$
If $(x, w) \notin R_L$, output $\perp$	If $(x, w) \notin R_L$, output $\perp$
Else output $(\text{PP}, \text{SP}, x, \pi, \text{STATE}_A)$	Else output $(\text{PP}', \text{SP}', x, \pi', \text{STATE}_A)$

   This is the "ADAPTIVE ZERO-KNOWLEDGE" requirement from [17, Definition 5], with the additional information SP resp. SP′ for A.

If those properties hold for every polynomial q, then we say that $(\mathcal{D}, P, V)$ enjoys *poly-adaptive security*.

Note that B is not computationally restricted. Also, we deviate slightly from the notation in [17], since with [17, Definition 5], it is not completely clear how x is chosen in the (non-adaptive) soundness requirement and how $|x'|$ relates to $|x|$.

We make two claims: first, for any q, our small modification makes the construction from [17] q-adaptively secure (in contrast to the original construction, which is *not* k-adaptively secure as argued above). Second, any q-adaptively secure DV-NIZK can be used to transform IND-CPA secure encryption into IND-q-CCA secure encryption (or even IND-CCA secure encryption in the case of $q =$"poly") using the construction from [15, 7, 17].

This second statement leaves of course open the question whether there is a (possibly even black-box) transformation from IND-CPA secure encryption to poly-adaptively secure DV-NIZK proofs. Namely, this would imply a transformation from IND-CPA secure encryption to IND-CCA secure encryption; this would not contradict [9], since the implied construction is in any case non-black box (since the techniques from [15, 7, 17] are), whereas [9] treats the case of black-box constructions.

Claim 5.3 Let $q = q(k)$ be a polynomial. Then the scheme $(\mathcal{D}, P, V)$ from Section 5.2 with $\ell := q(k) + k$ is a designated verifier non-interactive zero-knowledge proof system in the sense of [17, Definition 5] that is q-adaptively secure in the sense of Definition 5.2.

Proof: (Sketch.) The DV-NIZK property (including strong adaptive zero-knowledge) is standard and carries over from the scheme in [17] (which only uses $\ell = k$ instead of $\ell = q(k) + k$).

As for the q-adaptive soundness requirement, first note that the underlying NIZK proof system (P_1, V_1, P_2, V_2) used in $(\mathcal{D}, P, V)$ (i.e., Blum's hamiltonicity protocol) does not allow for tuples

(x, a, c_0, c_1) such that $x \notin L$ and *both* $(a, 0, c_0)$ and $(a, 1, c_1)$ are accepting transcripts (in the sense $V_2(x, a, b, c_b) = 1$ for $b = 0, 1$). Such tuples simply don't exist.

So to succeed in the q-adaptive soundness experiment, an adversary B must produce an $x \notin L$ and i (encrypted) tuples $(a_i, c_{0,i}, c_{1,i})$ with the following property. Namely, for all $i \in \{1, \ldots, \ell\}$, there is exactly one b_i with $V_2(x, a_i, b_i, c_{b,i}) = 1$, and it holds that $b_i = f_i$ for the f_i contained in the secret key SP. So information-theoretically, $f = (f_1, \ldots, f_\ell) \in \{0,1\}^\ell$ can be extracted from a successful DV-NIZK proof forger B. However, the public key SP is independent of f, and B may get at most $q = \ell - k$ bits of information about f from its oracle $\mathcal{O}^q_{\mathrm{PP,SP}}$. Information theoretic arguments show that hence, no B can "guess" (in an information-theoretic sense) f with non-negligible probability, and thus any given B must be unsuccessful in the q-adaptive soundness experiment. ▌

The next claim shows how to use the DV-NIZK to obtain IND-q-CCA security. This uses the construction of [17] that is in turn based on the construction of [7]. This latter construction was already used to achieve non-malleability even under chosen-ciphertext attacks (NM-CCA), which implies IND-CCA. In fact, the modification below achieves non-malleability in a setting where the adversary may, before generating the final forged ciphertext vector, ask for a limited number of decryption queries. Let's call this notion NM-q-CCA (a formal definition is obtained by equipping the adversary in the NME_b experiment from [17] with a limited decryption oracle DEC as in Definition 3.2). With the standard reduction, NM-q-CCA security is seen to imply IND-q-CCA security.

Claim 5.4 Let q be a polynomial. Let $(\mathcal{D}, P, V)$ be a designated verifier non-interactive zero-knowledge proof system in the sense of [17, Definition 5] that is q-adaptively secure in the sense of Definition 5.2. Then, $(\mathcal{D}, P, V)$ can be used in the construction of [17] to achieve NM-q-CCA (and thus IND-q-CCA) security. If $q =$"poly", then even NM-CCA (and thus IND-CCA) security is achieved.

Proof: (Sketch.) First, recall the scheme PKE scheme $\mathcal{PKE}$ constructed in [17] from $(\mathcal{D}, P, V)$ and an IND-CPA secure PKE scheme $\mathcal{pke}$. In $\mathcal{PKE}$, encryptions of m are of the form $(\mathbf{c}, \pi, \text{VKSIG}, \sigma)$, where $\mathbf{c} = (c_1, \ldots, c_k)$ is a vector of $\mathcal{pke}$ encryptions of m, 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 VKSIG.

Just as in [7], this whole construction is already geared towards NM-CCA security, and it fails to achieve full CCA security in the case $\ell = k$ considered in [17] only because of the weak soundness property of the employed DV-NIZK scheme. In fact, we only discuss the changes to the proof of [17] necessary to take care of the additional decryption oracle available to a q-CCA adversary attacking $\mathcal{PKE}$.

The games NME_b and $\mathsf{NME}_b^{(i)}$ (with $b = 0, 1$ and $i = 1, 2$) now allow the adversary to ask decryption queries in an adaptive manner. In these games of course, this can be done since the experiment itself knows all the secret keys, including the DV-NIZK key SP that is required to check the validity of a DV-NIZK proof. (In the case of the $\mathsf{NME}_b^{(2)}$ games, decryption of CCA decryption queries must be performed just like the decryption of the final forged ciphertext.) However, during some reductions which are performed when relating these games, an adversary A in the, say, NME_b experiment is mapped to a, say, adversary B on the zero-knowledge property of the DV-NIZK. In all cases relevant here, such a B internally simulates A and needs to simulate a NME_b setting for A. We will now go through the necessary changes to enable B to answer the additional decryption queries made by A.

First, in Claim 1 of [17], the experiments NME_b and $\mathsf{NME}_b^{(1)}$ (for $b = 0, 1$ respectively) are shown equivalent by a reduction on the adaptive zero-knowledge property of the DV-NIZK. Specifically, a theorem-chooser/distinguisher pair $(A_{\mathsf{zk}}, D_{\mathsf{zk}})$ on the DV-NIZK is constructed such that A_{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. A_{zk} generates all encryption and signature keypairs on its own, but takes the DV-NIZK public key PP from the adaptive zero-knowledge experiment. Thus, A_{zk} is generally *not* able to decrypt $\mathcal{PKE}$ ciphertexts since it cannot check the validity of the DV-NIZK part π. (Note that D_{zk} *is* able to decrypt since it is supplied with the corresponding DV-NIZK secret key SP by the zero-knowledge experiment.)

However, since we assume a DV-NIZK with a *strong* adaptive zero-knowledge property, in the corresponding reduction already A_{zk} knows SP and can thus answer decryption queries already before the challenge ciphertext is known. This is the only difference in the proof of Claim 1.

In Claim 2, the probability for the event $\mathtt{badNIZK}(\mathsf{Expt})$ 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 A_s on the soundness property of the DV-NIZK. Here, 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 [17] CPA setting, no decryptions are necessary, this is sufficient. However, in our q-CCA setting, A_s might need to answer up to q decryption queries in the NME_b 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 $\mathcal{O}^q_{\mathrm{PP,SP}}$.

Then, $\Pr\left[\mathsf{NME}_b^{(1)}\right] \approx \Pr[\mathsf{NME}_b]$, follows similarly (only now by a reduction on the *strong* adaptive zero-knowledge property as before). Now we cannot show $\Pr\left[\mathsf{NME}_b^{(1)}\right] = \Pr\left[\mathsf{NME}_b^{(2)}\right]$ (as in [17]), but we *can* show $\Pr\left[\mathsf{NME}_b^{(1)}\right] \approx \Pr\left[\mathsf{NME}_b^{(2)}\right]$, which sufficient for the further argument. The reason that we cannot show equality is that the view of an adversary in the $\Pr\left[\mathsf{NME}_b^{(i)}\right]$ 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 $\mathsf{NME}_b^{(2)}$, decryption is performed differently than in $\mathsf{NME}_b^{(1)}$). However, such decryption queries are answered differently only if event `badNIZK` or event `badSig` (which indicates that the adversary forged a signature) happens. The probability that one of these events occurs in $\mathsf{NME}_b^{(1)}$ is negligible, and thus $\Pr\left[\mathsf{NME}_b^{(1)}\right] \approx \Pr\left[\mathsf{NME}_b^{(2)}\right]$ follows.

In Claim 3 (which completes the proof in [17]), no properties of the DV-NIZK are used. ▌

Remark 5.5 With a small modification to the construction, one can do without the *strong* adaptive zero-knowledge property of the DV-NIZK (and instead only rely on the "ordinary," adaptive zero-knowledge property from [17] where the theorem-chooser does not get the secret key SP). Namely, one can use $\mathcal{PKE}$ as an IND-q-CCA secure *key encapsulation mechanism* (KEM) that on its own chooses a random message to encrypt instead of an IND-q-CCA secure PKE scheme. In this setting, a challenge ciphertext can be generated at the start of the experiment, and the need for the strong adapative zero-knowledge property in the changes to Claim 1 vanishes. Combined with an IND-CCA data encapsulation mechanism (which can be constructed from any one-way function), we obtain an IND-q-CCA secure PKE scheme.

Acknowledgements

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

References

[1] Ran Canetti, Yevgeniy Dodis, Shai Halevi, Eyal Kushilevitz, and Amit Sahai. Exposure-resilient functions and all-or-nothing transforms. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 453–469, Bruges, Belgium, May 14–18, 2000. Springer-Verlag, Berlin, Germany. (Cited on page 6.)

[2] Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 207–222, Interlaken, Switzerland, May 2–6, 2004. Springer-Verlag, Berlin, Germany. (Cited on page 2, 5.)

[3] Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Hugo Krawczyk, editor, *Advances in Cryptology – CRYPTO'98*, volume 1462 of *Lecture Notes in Computer Science*, pages 13–25, Santa Barbara, CA, USA, August 23–27, 1998. Springer-Verlag, Berlin, Germany. (Cited on page 2.)

[4] Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 45–64, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer-Verlag, Berlin, Germany. (Cited on page 2.)

[5] Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM Journal on Computing*, 33(1):167–226, 2003. (Cited on page 2.)

[6] Yevgeniy Dodis, Jonathan Katz, Shouhuai Xu, and Moti Yung. Key-insulated public key cryptosystems. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 65–82, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer-Verlag, Berlin, Germany. (Cited on page 2, 5.)

[7] Danny Dolev, Cynthia Dwork, and Moni Naor. Nonmalleable cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000. (Cited on page 1, 3, 8, 9, 10, 11.)

[8] 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. (Cited on page 5.)

[9] Yael Gertner, Tal Malkin, and Steven Myers. Towards a separation of semantic and cca security for public key encryption. In *Proceedings of TCC 2007*, pages ???–???, 2007. (Cited on page 2, 3, 10.)

[10] Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004. (Cited on page 2, 5.)

[11] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984. (Cited on page 4.)

[12] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st Annual ACM Symposium on Theory of Computing*, pages 44–61, Seattle, Washington, USA, May 15–17, 1989. ACM Press. (Cited on page 2.)

[13] L. Lamport. Constructing digital signatures from a one-way function. Technical Report CSL-98, SRI International, october 1979. (Cited on page 2, 5.)

[14] Y. Lindell. A simpler construction of cca2-secure public-key encryption under general assumptions. *Journal of Cryptology*, 19(3):359–377, 2006. (Cited on page 1, 3.)

[15] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd Annual ACM Symposium on Theory of Computing*, Baltimore, Maryland, USA, May 14–16, 1990. ACM Press. (Cited on page 1, 8, 9, 10.)

[16] R. Pass, A. Shelat, and V. Vaikuntanathan. Bounded-cca secure non-malleable encryption. manuscript, 2006. (Cited on page 3.)

[17] R. Pass, A. Shelat, and V. Vaikuntanathan. Construction of a non-malleable encryption scheme from any semantically secure one. In *Proceedings of CRYPTO 2006*, page ???, 2006. (Cited on page 2, 3, 8, 9, 10, 11, 12.)

[18] Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In Joan Feigenbaum, editor, *Advances in Cryptology – CRYPTO'91*, volume 576 of *Lecture Notes in Computer Science*, pages 433–444, Santa Barbara, CA, USA, August 11–15, 1992. Springer-Verlag, Berlin, Germany. (Cited on page 4.)

[19] John Rompel. One-way functions are necessary and sufficient for secure signatures. In *22nd Annual ACM Symposium on Theory of Computing*, pages 387–394, Baltimore, Maryland, USA, May 14–16, 1990. ACM Press. (Cited on page 2, 5.)

[20] Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *Proceedings of 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 1999. (Cited on page 1, 3.)