Theoretical Computer Science 572 (2015) 25–49

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Post-challenge leakage in public-key encryption ☆

Zongyang Zhang [a], Sherman S.M. Chow [b], Zhenfu Cao [c,*]

[a] Research Institute for Secure Systems, National Institute of Advanced Industrial Science and Technology (AIST), Japan
[b] Department of Information Engineering, Chinese University of Hong Kong, Hong Kong
[c] Shanghai Key Lab for Trustworthy Computing, East China Normal University, Shanghai, China

ARTICLE INFO

Article history:
Received 21 March 2013
Received in revised form 12 November 2014
Accepted 13 January 2015
Available online 16 January 2015
Communicated by G. Persiano

Keywords:
Public-key encryption
Chosen-ciphertext security
Identity-based encryption
Post-challenge leakage

ABSTRACT

When an adversary can measure the physical memory storing the decryption key, decryption functionality often comes in handy. Halevi and Lin (TCC'11) studied after-the-fact (or post-challenge) leakage in public-key encryption (PKE), in which an adversary can make leakage queries from a split state after seeing the challenge ciphertext, but left security against chosen-ciphertext attacks (CCA) as a future work. In this paper, we follow their work and formulate the definition of entropic leakage-resilient CCA-secure PKE, which we show can be realized by the Naor–Yung "double encryption" paradigm (STOC'90). We then leverage it to get a CCA-secure key-encapsulation mechanism in the presence of post-challenge leakage, in the same model of bounded memory leakage from a split state. Finally, we prove that the hybrid encryption framework is still applicable by presenting a construction of CCA-secure PKE in the presence of post-challenge leakage. As additional results, we extend these concepts to the identity-based setting, where many identity-based secret-keys can be leaked after the adversary got the challenge, and give a construction of identity-based encryption in the presence of post-challenge leakage in the split-state model, which can be instantiated by the identity-based hash proof systems of Alwen et al. (Eurocrypt'10) and Chow et al. (CCS'10).

1. Introduction

Traditional cryptographic schemes assume that an adversary has no control over the secret key. However, in real life, an attacker might gain some partial information about these secrets, by observing behavior of the protocol executions, or measuring the places where they are stored, via cold-boot attacks, electromagnetic measurements and many more. Several countermeasures have been proposed (e.g., tamper-proof hardware, oblivious RAM, etc.), but they are rather usually attack-specific and costly to implement.

Leakage-resilient cryptography [2] is introduced to formalize the above key-leakage attacks, and several models are proposed. In this paper, we consider the relative leakage model [3,4], where there is a natural bound on the leakage as a relative fraction of a secret key (e.g., less than half). Moreover, an adversary can repeatedly and adaptively learn arbitrary efficiently computable functions of the secret key, as long as the whole leakage amount is below the relative leakage bound.

☆ An extended abstract appeared in SCIS 2013, which has 9 pages in double-column form [1]. This is the full version which includes all formal definitions and security proofs, and the result on identity-based encryption.

* Corresponding author.
E-mail addresses: zongyang.zhang@gmail.com (Z. Zhang), sherman@ie.cuhk.edu.hk (S.S.M. Chow), zfcao@sei.ecnu.edu.cn (Z. Cao).

http://dx.doi.org/10.1016/j.tcs.2015.01.010
0304-3975/

There are already many results on leakage-resilient public-key encryption (PKE) and identity-based encryptions (IBE) [5–15]. However, security models of these works (except Halevi and Lin's result [12]) have a constraint that no leakage query is allowed after the challenge ciphertext is given. Recall that the semantic security of PKE (either against chosen-plaintext attacks (CPA) or chosen-ciphertext attacks (CCA)[1]) requires no probabilistic polynomial time (PPT) adversary, submitting two equal-length messages m_0, m_1 and receiving a (challenge) ciphertext, can tell if it is encrypting m_0 or m_1. If we allow an adversary to get key leakage after receiving the challenge ciphertext, the adversary can simply ask for a function that decrypts the challenge and outputs a single bit b when the decryption result is m_b. Thus, it is impossible to achieve the semantic security in standard security models.

This inherent constraint heavily affects the meaning and hence the applicability of previous leakage-resilient PKE. Even only a bit of a secret key is leaked, there is no guarantee regarding the security of all previously encrypted messages. Halevi and Lin [12] gave an example scenario about disk encryption to illustrate that it is reasonable to assume an attacker can see the encrypted (disk) content before measuring the memory (and getting the leakage of the secret key). Towards resolving this problem, they studied *after-the-fact leakage* (or called post-challenge leakage), in which the adversary is able to obtain leakage information after seeing the challenge ciphertext. They first formulated the notion of entropic leakage-resilient PKE, and gave a construction meeting it. This notion captures the intuition that even if an adversary learns t-bit information about the secret key, it cannot learn more than t bits about the plaintext. They then leveraged the notion of entropic leakage-resilient PKE to give a construction of leakage-resilient PKE in the presence of post-challenge leakage, in a *split-state model*, where the key is broken into parts, and the adversary is only allowed to get leakage from every part separately but not a global leakage from the entire secret key. To fit with this model, one may store two parts of the key in two independent computer chips, which is a moderate cost for achieving post-challenge leakage-resilience, which was not possible otherwise.

Halevi and Lin only considered chosen-plaintext attacks in the presence of post-challenge leakage. For many scenarios including the aforementioned disk encryption application, it is reasonable to assume that an adversary can mount chosen-ciphertext attacks [16] when leakage-attack is possible. Note that since the leakages are bounded and they come from different parts of a secret key separately, a general decryption functionality is not available to the adversary. Indeed, Halevi and Lin remarked that "it will be very interesting" to explore other notions of security in the context of post-challenge leakage, for instance, CCA-secure encryption.

1.1. Our contribution

Public-key encryption In this paper, we address this problem by constructing the first CCA-secure PKE in the presence of post-challenge leakage, in a split-state model. Toward doing this, we first extend the notion of entropic leakage-resilient CPA-secure PKE to entropic leakage-resilient CCA-secure one which also supports labels.[2] This notion considers the entropy of the encrypted message under chosen-ciphertext attacks and in the presence of post-challenge leakage. Roughly speaking, even if an adversary can adaptively make decryption queries and choose leakage function on the secret key for learning partial information about it, nothing can be learned more than what has been leaked on the plaintext.

Next we show that by applying the Naor–Yung "double encryption" [17] paradigm, one can combine any entropic leakage-resilient CPA-secure PKE with any (ordinary) CPA-secure encryption scheme, together with a one-time simulation-sound non-interactive zero-knowledge (NIZK) proof supporting labels, to obtain an entropic leakage-resilient CCA-secure labeled-PKE.

In order to get a leakage-resilient CCA-secure PKE in the split-state model, we first design a leakage-resilient CCA-secure key-encapsulation mechanism (KEM). We use two instances of an entropic leakage-resilient CCA-secure PKE supporting labels, and a strong one-time signature scheme secure against one-time chosen message attack. To get an encapsulated symmetric key, we choose two random strings x_0, x_1, encrypt each x_j in ciphertext c_j under a different copy of the entropic leakage-resilient PKE with vk as the label, where (vk, sk) is a signature key pair we freshly generated, then we generate a signature σ on the two ciphertexts c_0, c_1, and use a two-source extractor to extract the symmetric key $k = \mathsf{Ext}_2(x_0, x_1)$. To decrypt we first check the validity of the signature, then retrieve the strings x_0, x_1 by decrypting the corresponding ciphertext, and recover the symmetric key k. The main reason that this approach works is as follows. As long as an adversary cannot choose leakage function on both parts of the secret key simultaneously, and the decryption queries do not leak information on the encrypted challenge message (due to a similar design which achieves CCA-security of multiple-encryption [18]), the underlying entropic leakage-resilient CCA-secure PKE guarantees that x_0 and x_1 still have many entropy left, and the two-source extractor Ext_2 still extracts out the randomness.

Finally, following the hybrid encryption paradigm, we combine a leakage-resilient CCA-secure KEM in the split-state model with a (one-time) symmetric key encryption to obtain a leakage-resilient CCA-secure PKE in the split-state model.

Identity-based encryption With the wide applications of identity-based encryption (IBE) in mind, we also develop a notion of entropic leakage-resilient IBE, in which the min-entropy of a random message is still "high" given a ciphertext of this

[1] By CCA we mean a-posteriori chosen-ciphertext attacks, which are usually called CCA2 instead.

[2] In many applications where one uses a CCA-secure encryption scheme, the notion of a label is very useful. Very briefly, a label consists of public data which is non-malleably attached to a ciphertext.