Advances in Cryptology – Proceedings of ASIACRYPT '99 (14 – 18 november 1999, Singapore)
K. Y. Lam, E. Okamoto and C; Xing Eds. Springer-Verlag, LNCS 1716, pages 165–179.

Efficient Public-Key Cryptosystems Provably Secure Against Active Adversaries

Pascal Paillier[1] and David Pointcheval[2]

[1] Gemplus Cryptography Department, 34 Rue Guynemer, 92447 Issy-Les-Moulineaux, France and ENST, 46 Rue Barrault, 75634 Paris Cedex 13, France. Pascal.Paillier@gemplus.com
[2] LIENS – CNRS, École Normale Supérieure, 45 rue d'Ulm, 75230 Paris Cedex 05, France. David.Pointcheval@ens.fr, http://www.dmi.ens.fr/∼pointche.

Abstract. This paper proposes two new public-key cryptosystems semantically secure against *adaptive chosen-ciphertext attacks*. Inspired from a recently discovered trapdoor technique based on composite-degree residues, our converted encryption schemes are proven, in the random oracle model, secure against active adversaries (NM-CCA2) under the assumptions that the Decision Composite Residuosity and Decision Partial Discrete Logarithms problems are intractable. We make use of specific techniques that differ from Bellare-Rogaway or Fujisaki-Okamoto conversion methods. Our second scheme is specifically designed to be efficient for decryption and could provide an elegant alternative to OAEP.

1 Introduction

Diffie and Hellman's famous paper [7] initiated the paradigm of asymmetric cryptography in the late seventies but since, very few trapdoor mechanisms were found that fulfill satisfactory security properties. Of course, the first security criterion a cryptosystem has to verify is the one-wayness of its encryption function, but this notion does not suffice to evaluate (and get people convinced of) the strength of an encryption scheme.

A typical example is RSA [18] which, although very popular and widely used in many cryptographic applications, suffers from being *malleable* and consequently requires an additional treatment (some probabilistic padding) on the plaintext in order to strengthen its practical security. Resistance against chosen-ciphertext attacks, in this case, relies on the conjoint use of an external paradigm instead of being inherently provided, although this empirical approach may sometimes appear insufficient, as shown by Bleichenbacher [4] and more recently by Coron, Naccache and Stern [5]. This motivates the construction of provably secure padding techniques such as OAEP [3] or Fujisaki-Okamoto [10].

Considerable efforts have recently been made to investigate cryptosystems achieving provable security against active adversaries at reasonable encryption and/or decryption cost. Our paper introduces two such cryptosystems that are efficient for decryption and meet provable security at the strongest level (NM-CCA2) in the random oracle model. We make use of specific techniques that differ from those of [3] and [10].

We begin by briefly surveying known notions of security for public-key encryption schemes, refering the reader to [1] for their formal definitions and connections.

1.1 Notions of Security

Formalizing another security criterion that an encryption scheme should verify beyond one-wayness, Goldwasser and Micali [11] introduced the notion of *semantic security*. Also called *indistinguishability of encryptions* (or IND for short), this property captures the idea according to which an adversary should not be able to learn any information whatsoever about a plaintext, its length excepted, given its encryption. The property of *non-malleability* (NM), independently proposed by Dolev, Dwork and Naor [8], supposes that, given the encryption of a plaintext x, the attacker cannot produce the encryption of a related plaintext x'. Here, rather than learning some information about x, the adversary will try to output the encryption of x'. These two properties are related in the sense that non-malleability implies semantic security for any adversary model, as pointed out in [8] and [1].

On the other hand, there exist several types of adversaries, or attack models. In a chosen-plaintext attack (CPA), the adversary has access to an encryption oracle, hence to the encryption of any plaintext she wants. Clearly, in a public-key setting, this scenario cannot be avoided. Naor and Yung [13] consider non-adaptive chosen-ciphertext attacks (CCA1) (also known as lunchtime or midnight attacks), wherein the adversary gets, in addition, access to a decryption oracle before being given the challenge ciphertext. Finally, Rackoff and Simon [17] defined adaptive chosen-ciphertext attacks (CCA2) as a scenario in which the adversary queries the decryption oracle before and *after* being challenged; her only restriction here is that she may not feed the oracle with the challenge ciphertext itself. This is the strongest known attack scenario.

Various security levels are then defined by pairing each goal (IND or NM) with an attack model (CPA, CCA1 or CCA2), these two caracteristics being considered separately. Interestingly, it has been shown that IND-CCA2 and NM-CCA2 were strictly equivalent notions [1].

Beyond this, Bellare and Rogaway [3] proposed the concept of *plaintext awareness*, where the adversary attempts to produce a valid ciphertext without knowing the corresponding plaintext. This additional security notion was only properly defined in the random oracle model [2].

1.2 The Random Oracle Model

The random oracle model was proposed by Bellare and Rogaway [2] to provide heuristic (yet satisfactorily convincing) proofs of security. In this model, hash functions are considered to be ideal, *i.e.* perfectly random. From a security viewpoint, this impacts all three adversary models by giving the attacker an additional access to the random oracles of the scheme.

1.3 Related Work

The basic El Gamal encryption scheme [9], which one-wayness relates to the celebrated Diffie-Hellman (DH) problem, was recently proven semantically secure (*i.e.* secure in the sense of IND-CPA) by Tsiounis and Yung [22] under

the Decision Diffie-Hellman (D-DH) assumption. However, just like RSA, the original scheme remains totally unsecure regarding active attacks. The same authors therefore proposed a converted scheme provably secure in the sense of NM-CCA2 in the random oracle model, under the D-DH assumption in addition to a non-standard one. Independently, Shoup and Gennaro [20] proposed another converted scheme NM-CCA2 in the random oracle model under the D-DH assumption only. The same year, Cramer and Shoup [6] also presented an El Gamal-based cryptosystem, the first to be simultaneously pratical and provably NM-CCA2 secure in the standard model, provided that the D-DH assumption holds.

Several authors have investigated other intractability assumptions. Pointcheval [16] proposed DRSA, an encryption scheme based on the Dependent-RSA Problem, and provided efficient variants provably NM-CCA2 secure in the random oracle model under the hypothesis that the decisional version of the D-RSA Problem is intractable. Naccache and Stern [12], and independently Okamoto and Uchiyama [14] investigated different approaches based on high degree residues. The one-wayness (resp. semantic security) of their schemes is ensured by the Prime Residuosity assumption (resp. the hardness of distinguishing prime-degree residues). Finally, Paillier [15] proposed an encryption scheme based on composite-degree residues wherein semantic security relies on a similar assumption (see below).

In 94, Bellare and Rogaway [3] proposed OAEP, a specific hash-based treatment applicable to any one-way trapdoor permutation to make it secure in the sense of NM-CCA2. Standing in the random oracle model, their security proof is widely recognized and initiated the upcoming RSA-based PKCS #1 V2.0 standard [19]. More recently, Fujisaki and Okamoto [10] discovered a generic conversion method which transforms any semantically secure encryption scheme into a scheme secure in the sense of NM-CCA2 in the random oracle model. The conversion is low-cost for encryption (one additional hash), but appears to be heavy for decryption[1].

1.4 Outline of the Paper

In this paper, we propose two new encryption schemes that are provably secure against adaptive chosen-ciphertext attacks (NM-CCA2) in the random oracle model. Based on Paillier's probabilistic encryption schemes [15], we provide semantic security relatively to two number-theoretic decisional problems, namely the Decision Composite Residuosity and Decision Partial Discrete Log problems. With an efficiency comparable to OAEP for decryption, we believe that the second of these cryptosystems could provide an elegant alternative to the new standard.

2 The Basic Schemes

This section briefly describes the public-key cryptosystems proposed in [15], keeping the same notations as in the original paper.

[1] the converted decryption process includes a complete data re-encryption.

2.1 Notations

We set $n = pq$ where p and q are large primes. We will denote by ϕ Euler's totient function and by λ Carmichael's function on n, *i.e.* $\phi = (p-1)(q-1)$ and $\lambda = \mathrm{lcm}(p-1, q-1)$ in the present case. For technical reasons, we will focus on moduli $n = pq$ such that $\gcd(p-1, q-1) = 2$, which yields $\phi = 2\lambda$. Recall that $|\mathbb{Z}_{n^2}^*| = \phi(n^2) = n\phi$ and that Carmichael's theorem implies that

$$\forall w \in \mathbb{Z}_{n^2}^* , \quad \begin{cases} w^{\lambda} = 1 \bmod n \\ w^{n\lambda} = 1 \bmod n^2 \end{cases} ,$$

We denote by RSA $[n, e]$ the well-known problem of extracting e-th roots modulo n where $n = pq$ is of unknown factorization.

2.2 Setting

Let $n = pq$ be a modulus chosen as above and $g \in \mathbb{Z}_{n^2}^*$. It is known that the integer-valued function $\mathcal{E}_g$ defined as

$$\begin{array}{rcl} \mathbb{Z}_n \times \mathbb{Z}_n^* & \longmapsto & \mathbb{Z}_{n^2}^* \\ (x, y) & \longmapsto & g^x \cdot y^n \bmod n^2 \end{array}$$

is a bijection if the order of g in $\mathbb{Z}_{n^2}^*$ is a multiple of n. When this condition is met, then given $w \in \mathbb{Z}_{n^2}^*$, the unique integer x for which there exists a y such that $\mathcal{E}_g(x, y) = w$ is called the (n-residuosity) class of w and is denoted $[\![w]\!]_g$. It is believed that for given n, g and w the problem of computing the class $[\![w]\!]_g$ of w is computationally hard: this is known as the composite residuosity assumption (CRA) [15].

It has been shown, however, that the knowledge of the factors p and q is sufficient for computing the class of any integer $w \in \mathbb{Z}_{n^2}^*$. Indeed, setting

$$\mathcal{S}_n = \{u < n^2 \mid u = 1 \bmod n\} ,$$

and

$$\forall u \in \mathcal{S}_n \quad \mathrm{L}(u) = (u-1)/n ,$$

we have

$$[\![w]\!]_g = \frac{\mathrm{L}(w^{\lambda} \bmod n^2)}{\mathrm{L}(g^{\lambda} \bmod n^2)} \bmod n . \tag{1}$$

2.3 Description

First, randomly select an integer g such that n divides the order of g. This can be done by checking whether

$$\gcd\left(\mathrm{L}(g^{\lambda} \bmod n^2), n\right) = 1 . \tag{2}$$

The pair (n, g) is then published as the public key, whilst the pair (p, q) (or equivalently λ) forms the secret key. The cryptosystem is described on figure 1.

Decryption thus requires essentially one exponentiation modulo n^2 with exponent λ. As pointed out in [15], this encryption scheme is one-way if and only if the CRA holds.

Initialization

$n = pq$, $g \in \mathbb{Z}^*_{n^2}$
n divides the order of g

PUBLIC: n, g
SECRET: λ

Encryption

plaintext: $m \in \mathbb{Z}_n$
select a random $r \in \mathbb{Z}^*_n$
ciphertext: $c = g^m \cdot r^n \bmod n^2$

Decryption

ciphertext: $c \in \mathbb{Z}^*_{n^2}$
plaintext: $m = \dfrac{\mathrm{L}(c^\lambda \bmod n^2)}{\mathrm{L}(g^\lambda \bmod n^2)} \bmod n$

Fig. 1. Main Scheme

2.4 The Subgroup Variant

In this variant (see figure 2), the idea consists in restricting the ciphertext space $\mathbb{Z}^*_{n^2}$ to the subgroup $\langle g \rangle$ of smaller order by taking advantage of the following extension of Equation 1. Assume that the order of g is $n\alpha$ for some $1 \leq \alpha \leq \lambda$. Then for any $w \in \langle g \rangle$,

$$[\![w]\!]_g = \frac{\mathrm{L}(w^\alpha \bmod n^2)}{\mathrm{L}(g^\alpha \bmod n^2)} \bmod n \ . \tag{3}$$

By carefully setting α to an integer of suitable length ℓ (typically 320 bits in practice), the decryption workload thus decreases to an exponentiation with an ℓ-bit exponent.

Initialization

$n = pq$, $\alpha | \lambda$
$h \in \mathbb{Z}^*_{n^2}$ of maximal order $n\lambda$
$g = h^{\lambda/\alpha} \bmod n^2$

PUBLIC: n, g
SECRET: α

Encryption

plaintext: $m \in \mathbb{Z}_n$
randomly select $r < 2^\ell$
ciphertext: $c = g^{m+nr} \bmod n^2$

Decryption

ciphertext: $c \in \mathbb{Z}^*_{n^2}$
plaintext: $m = \dfrac{\mathrm{L}(c^\alpha \bmod n^2)}{\mathrm{L}(g^\alpha \bmod n^2)} \bmod n$
if the computation was impossible, output "failure"

Fig. 2. Subgroup Variant

Naturally, the problem of computing $[\![w]\!]_g$ for $w \in \langle g \rangle$ is (computationally) weaker than doing so for $w \in \mathbb{Z}_{n^2}^*$. For $\ell = \Omega(|n|^\epsilon)$ with $\epsilon > 0$, it is however considered to be intractable: this complexity hypothesis is known as the Partial Discrete Log (PDL) assumption. Moreover, inverting the hereabove encryption scheme was shown to be intractable if (and only if) the PDLA assumption holds.

2.5 Security Results

In a similar way, both cryptosystems were proven semantically secure against chosen-plaintext attacks (IND-CPA) under the additional complexity assumptions that the decisional versions of the Composite Residuosity and Partial Discrete Log problems are also intractable (see [15] for technical details). These intractability hypothesis are called Decision Composite Residuosity (D-CRA) and Decision Partial Discrete Log (D-PDLA) assumptions, respectively. Because of their obvious malleability, however, both cryptosystems do not resist chosen-ciphertext attacks.

Remark 1. Interestingly, adaptive attacks do not seem to allow a total breakdown (secret key retrieval) of these encryption schemes, whilst they trivially do in Okamoto-Uchiyama's [14] as pointed out in their original paper.

In the next section, we show how to render these schemes secure against adaptive chosen-ciphertext attacks relatively to the D-CRA and the D-PDLA assumptions in the random oracle model.

3 Improved Cryptosystems

The notion of security against an adaptive chosen-ciphertext attack (IND-CCA2) was introduced by Rackoff and Simon [17] as the property that a cryptosystem must have to resist active adversaries. In this scenario, the adversary makes queries of her choice to a decryption oracle during two stages. After the first stage of queries (the "find" stage), the attacker chooses two messages and requests an encryption oracle to encrypt one of them, leaving to the oracle the (secret) choice of which one. The adversary then continues to query the decryption oracle (the "guess" stage) with ciphertexts of her choice. Finally, she tells her guess about the choice the encryption oracle made. If she correctly guesses with probability non-negligibly higher than one half, in polynomial time, then the encryption scheme is considered unsecure.

In this section, we propose a modification of the main scheme (*c.f.* section 2.3) which provides security in the sense of IND-CCA2 under the D-CR assumption in the random oracle model. At that point, note that OAEP [3] cannot be employed *ad hoc* for this purpose, unless using Paillier's trapdoor one-way permutation: this would lead to a practical but inefficient converted encryption scheme. Also, Fujisaki and Okamoto's conversion method [10] could theoretically be applied but would considerably decrease the decryption speed (as previously said, a re-encryption is necessary during the converted decryption process) and therefore significatively reduce the practical interest of the subgroup variant.

Instead of these approaches, we rather rely on modifications specifically adapted to the schemes. For this purpose, we use a t-bit random number and two hash functions $G, H : \{0,1\}^* \mapsto \{0,1\}^{|n|}$ seen as random oracles. The encryption scheme is described on figure 3.

Initialization

$n = pq$, $g \in \mathbb{Z}^*_{n^2}$
n divides the order of g

PUBLIC: n, g
SECRET: λ

Encryption

plaintext: $m < 2^{|n|-t-1}$
randomly select $r < 2^t$
$z = H(m,r)^n \bmod n^2$
$M = m||r + G(z \bmod n) \bmod n$
ciphertext: $c = g^M z \bmod n^2$

Decryption

ciphertext: $c \in \mathbb{Z}^*_{n^2}$
$M = \dfrac{\mathrm{L}(c^\lambda \bmod n^2)}{\mathrm{L}(g^\lambda \bmod n^2)} \bmod n$
$z' = g^{-M} c \bmod n$
$m'||r' = M - G(z') \bmod n$
if $H(m',r')^n = z' \bmod n$ then the plaintext is m'
otherwise output "failure"

Fig. 3. Encryption scheme secure against adaptive attacks (Scheme 1)

Theorem 2. *Provided $t = \Omega\left(|n|^\delta\right)$ for $\delta > 0$, Scheme 1 is semantically secure against adaptive chosen-ciphertext attacks under the Decision Composite Residuosity assumption in the random oracle model.*

Proof. Let us consider an adversary $A = (A_1, A_2)$ against the semantic security of Scheme 1, where A_1 denotes the "find"-stage and A_2 the "guess"-stage. We then use this adversary to efficiently decide n-residuosity classes. Indistinguishability of encryptions will be due to the randomness of the oracle G, whereas adaptive attacks are covered thanks to the random oracle H.

Simulation of the Decryption Oracle Since we consider chosen-ciphertext attacks, the adversary has access, in both stages, to a decryption oracle $\mathcal{D}$ that we have to simulate: when the attacker asks for a ciphertext c to be decrypted, the simulator checks in the query-answer history obtained from the random oracle H whether some entry leads to the ciphertext c and then returns m; otherwise, it returns "failure". This provides a quasi-perfect simulation since the probability of producing a valid ciphertext without asking the query (m, r) to the random oracle H (whose answer a has to satisfy the test $a^n = z \bmod n$) is upper-bounded

by $1/\phi(n) \leq 2/n$, which is clearly negligible. This simulator can also be seen as a knowledge extractor, which provides the plaintext-awareness [3] of the scheme.

Semantic Security We will rely on the attacker A to design a distinguisher B for n-residuosity classes. Let (w, α) be a given instance of the D-CR problem: α is suspected to be the n-residuosity class of w.

Our distinguisher B first randomly chooses $u \in \mathbb{Z}_n$, $v \in \mathbb{Z}_n^*$ and $0 \leq r < 2^t$ and then computes $z = w \cdot g^{-\alpha} v^n \bmod n$ as well as $c = w \cdot g^u v^n \bmod n^2$. It then runs A_1 and gets two messages m_0 and m_1. B chooses a bit b and runs A_2 on the ciphertext c, supposed to be the ciphertext of m_b using the random r.

If during the game z is asked to the oracle G (which event will be denoted by AskG), one stops the game and B returns 1. If (m_0, r) or (m_1, r) are asked to the oracle H (which event will be denoted by AskH), one stops the game and B returns 0. In any other case, B returns 0 when A_2 ends.

One has to remark that no more than one event AskG or AskH is likely to happen since both events make B terminate the game. Furthermore, because of the random choice of r, the probability of AskH is upper bounded by $q_H/2^t$ in any case, where q_H denotes the number of queries asked to H.

Since G and H are seen like random oracles, the attacker has no chance to correctly guess b, during a real attack (or in the case where $\alpha = [\![w]\!]_g$), if none of the events AskG or AskH occur, then $\mathsf{Adv}_A \leq \Pr[\mathsf{AskG} \vee \mathsf{AskH} | [\![w]\!]_g = \alpha]$.

On the other hand, if $\alpha \neq [\![w]\!]_g$, z is perfectly random (and furthermore independent of c), then AskG cannot occur with probability greater than $q_G/\phi(n)$, where q_G denotes the number of queries asked to G.

Therefore, the distinguisher B gets the following advantage in deciding the n-residuosity classes

$$\begin{aligned}
\mathsf{Adv}_B &= \Pr[1 | [\![w]\!]_g = \alpha] - \Pr[1 | [\![w]\!]_g \neq \alpha] \\
&= \Pr[\mathsf{AskG} | [\![w]\!]_g = \alpha] - \Pr[\mathsf{AskG} | [\![w]\!]_g \neq \alpha] \\
&= \Pr[\mathsf{AskG} \vee \mathsf{AskH} | [\![w]\!]_g = \alpha] - \Pr[\mathsf{AskH} | [\![w]\!]_g = \alpha] - \Pr[\mathsf{AskG} | [\![w]\!]_g \neq \alpha] \\
&\geq \mathsf{Adv}_A - q_H/2^t - q_G/\phi(n) \geq \mathsf{Adv}_A - q_H/2^t - 2q_G/n.
\end{aligned}$$

Reduction Cost To conclude, if there exists an active attacker A against semantic security, one can decide n-residuosity classes with an advantage greater than

$$\mathsf{Adv}_A \times \left(1 - \frac{2}{n}\right)^{q_D} - \frac{q_H}{2^t} - \frac{2q_G}{n} \geq \mathsf{Adv}_A - \frac{q_H}{2^t} - 2 \cdot \frac{q_G + q_D}{n},$$

where q_D, q_G and q_H denote the number of queries asked to the decryption oracle, G and H respectively. □

However, the ability to decide n-residues may not be enough to break the semantic security. Furthermore, one can also prove that the converted scheme is still one-way relatively to the computational problem.

Theorem 3. *Provided $t = \Omega\left(|n|^\delta\right)$ for $\delta > 0$, Scheme 1 is one-way, even against adaptive chosen-ciphertext attacks, under the Composite Residuosity assumption in the random oracle model.*

Proof. Let us consider an adversary A able to decrypt any ciphertext c with probability ε, within a time bound T. Since we consider chosen-ciphertext attacks, the adversary has access to a decryption oracle $\mathcal{D}$ whose simulation works as described above. During an attack, two cases may appear

case 1: either the attacker tries to check the validity of the ciphertext, then she has to compute $H(m, r)$ which gives us m, r and therefore the n-residuosity class of the ciphertext;

case 2: or she asks the query z to the oracle G (as previously noticed, she gets no information about the plaintext without such a query).

As already mentioned, the attacker cannot get any information about the plaintext if none of these cases applies. Then either the first case applies with probability greater than $\varepsilon/2$, and therefore the attacker can be used to compute n-residuosity classes with probability greater than $\varepsilon/2$ within a time bound T, or the second case applies with probability greater than $\varepsilon/2$. Let us consider the second case. For simplicity, we will restrict the study to the setting $3t \leq |n|$ since other parameter choices present no practical interest, and would make the proof be more intricate.

Let w be an element of $\mathbb{Z}_{n^2}^*$ of class α we want to compute. One randomly chooses $\alpha_0, \alpha_1 \in \mathbb{Z}_n$, $\beta_0, \beta_1 \in \mathbb{Z}_n^*$, and computes

$$w_0 = w \cdot g^{\alpha_0} \beta_0^n \bmod n^2 \quad \text{and} \quad w_1 = w^{-2^t} \cdot g^{\alpha_1} \beta_1^n \bmod n^2.$$

The ciphertexts w_0 and w_1 are successively given to the attacker and all the answers ρ_i (randomly chosen in $\mathbb{Z}_n$ by the simulation of G during the attack against w_0) as well as all the answers σ_i (given during the attack against w_1) are collected and stored. Because of the uniform distribution of w_0 and w_1 in $\mathbb{Z}_{n^2}^*$, the attacker A succeeds in correctly finding m_0 and m_1, dropping in case 2, with probability $\varepsilon^2/4$. Then, there exist $0 \leq r_0, r_1 < 2^t$ and indices i, j such that

$$\begin{aligned} \alpha + \alpha_0 &= 2^t m_0 + r_0 + \rho_i \bmod n \\ -2^t \cdot \alpha + \alpha_1 &= 2^t m_1 + r_1 + \sigma_j \bmod n. \end{aligned}$$

By combination, one gets

$$2^t r_0 + r_1 = 2^t \alpha_0 + \alpha_1 - 2^{2t} m_0 - 2^t m_1 - 2^t \rho_i - \sigma_j \bmod n.$$

Hence there exists at least one pair (i, j) such that

$$0 \leq 2^t \alpha_0 + \alpha_1 - 2^{2t} m_0 - 2^t m_1 - 2^t \rho_i - \sigma_j \bmod n < 2^{2t}.$$

One randomly chooses such a pair (i, j) which allows to compute r_0 and r_1, and therefore α (with probability greater than $1/q_G^2$ if m_0 and m_1 are correct, since there are at most q_G^2 possible pairs (i, j), where q_G is the number of queries asked to G during a decryption). Consequently, our reduction recovers α with probability greater than $(\varepsilon/2q_G)^2$ within a time bound $2T$.

However, this reduction can be heuristically shown much more efficient. Indeed, the probability of having one valid pair (i, j) for incorrect plaintexts

(m_0, m_1) and the probability of having many pairs (i, j) for valid plaintexts m_0, m_1 are both upper-bounded by $q_G^2/2^t$ (because of the randomness of the ρ_k and σ_k sequences, and perfect independence of $\alpha_0, \alpha_1, \beta_0, \beta_1$). We can therefore ignore them for a large enough security parameter t. Finally, one can recover α within an expected time bounded by $8T/\varepsilon$ (this is an optimal reduction). □

At this point, we comment that Fujisaki and Okamoto's conversion technique [10] would have given an identical security level and a quite similar computational workload. The superiority of our approach resides in that

a) the one-wayness of our both converted schemes are equivalent to the CR and PDL problems, whereas Fujisaki-Okamoto would have inherently restricted one-wayness to the *decision* problems D-CR and D-PDL,

b) the same proof as above will now apply almost unchanged on the subgroup variant, leading to a far better decryption efficiency (our validity test does not involve a complete re-encryption).

We now turn to show how to modify the subgroup variant (*c.f.* section 2.4) to meet NM-CCA2 security under the D-PDLA in the random oracle model. As before, we make use of two hash functions, $G, H : \{0,1\}^* \mapsto \{0,1\}^{|n|}$ considered as random oracles. In what follows, we set α to an odd divisor $\alpha = 2a + 1$ of λ with bitsize ℓ and randomly pick an element h of maximal order $n\lambda$ in $\mathbb{Z}_{n^2}^*$. Recall that the modulus n is chosen such that $p-1$ and $q-1$ do not have common prime divisors other than 2. The encryption scheme is depicted on figure 4.

Initialization

$n = pq$, $\gcd(p-1, q-1) = 2$, $\alpha = 2a + 1 | \lambda$
$h \in \mathbb{Z}_{n^2}^*$ of maximal order $n\lambda$
$g = h^{\lambda/\alpha} \bmod n^2$

PUBLIC: n, g
SECRET: α

Encryption

plaintext: $m < 2^{|n|-t-1}$
randomly select $r < 2^t$
$z = g^{nH(m,r)} \bmod n^2$
$M = m||r + G(z \bmod n) \bmod n$
ciphertext: $c = g^M z \bmod n^2$

Decryption

ciphertext: $c \in \mathbb{Z}_{n^2}^*$

$$M = \frac{\mathrm{L}(c^\alpha \bmod n^2)}{\mathrm{L}(g^\alpha \bmod n^2)} \bmod n$$

if the computation was impossible, output "failure"
$z' = g^{-M} c \bmod n$
$m'||r' = M - G(z') \bmod n$
if $g^{nH(m',r')} = z' \bmod n$ then the plaintext is m'
otherwise output "failure"

Fig. 4. Efficient variant secure against adaptive attacks (Scheme 2)

Theorem 4. *Provided $t = \Omega\left(|n|^{\delta}\right)$ and $\ell = \Omega\left(|n|^{\epsilon}\right)$ for $\delta, \epsilon > 0$, Scheme 2 is semantically secure against adaptive chosen-ciphertext attacks under the D-PDLA in the random oracle model.*

Proof. The proof is essentially the same, using $g^{H(m,r)}$ instead of $H(m,r)$, but we also have to prove that one cannot decrypt a ciphertext which has not been correctly computed (using the encryption scheme). Indeed, our simulation of the decryption oracle (the plaintext extractor) can only decrypt a valid ciphertext.

Let then c be an accepted ciphertext (*i.e* which has not caused a decryption failure). This means that $c^{\alpha} = 1 \bmod n$ and thus $c^{\alpha} \in \mathcal{S}_n = \langle h^{\lambda} \rangle \subset \langle h \rangle$. We have assumed that $\alpha = 2a + 1$, $\gcd(p-1, q-1) = 2$ and h of maximal order $\lambda(n^2) = \phi(n^2)/2$. Then, for any $x \in \mathbb{Z}_{n^2}^*$, $x^2 \in \langle h \rangle$. Therefore,

$$c = c^{\alpha} \cdot c^{-2a} = c^{\alpha} \cdot (c^2)^{-a} \in \langle h \rangle \ ,$$

which implies the existence of an x such that $c = h^x \bmod n^2$. Furthermore, $c^{\alpha} = h^{x\alpha} = 1 \bmod n$, and h is also of maximal order in $\mathbb{Z}_n^*$. Therefore there exists y such that $x = \beta y$, where $\lambda = \alpha\beta$. Thus, $c = h^{\beta y} = g^y \bmod n^2$, and the n-residuosity class M obtained during the decryption process satisfies $M = y \bmod n$ and $c = g^{M+kn} \bmod n^2$. Hence $z = g^{kn} \bmod n$.

Because the ciphertext was accepted, we have that $z = g^{nH(m',r')} \bmod n$, and therefore $k = H(m', r') \bmod \alpha$, because g is of order α in $\mathbb{Z}_n^*$, and so is g^n (n is relatively prime to ϕ). Finally, if we define $m = m'$ and $r = r'$, one gets c as ciphertext. □

Theorem 5. *Provided $t = \Omega\left(|n|^{\delta}\right)$ and $\ell = \Omega\left(|n|^{\epsilon}\right)$ for $\delta, \epsilon > 0$, Scheme 2 is one-way, even against adaptive chosen-ciphertext attacks, under the Partial Discrete Logarithm assumption in the random oracle model.*

Remark 6. Because of the plaintext extractor presented in the proof of Theorem 2, both schemes are plaintext-aware [3].

4 Encryption Parameters

In practice, α should be typically set to a 320-bit divisor of λ such that $\alpha = \alpha_p \alpha_q$ where α_p divides $p-1$ but not $q-1$ and α_q divides $q-1$ but not $p-1$. This can be met using an appropriate key generation algorithm. Note that our converted schemes, like the original ones, allow Chinese remaindering for decryption. In the subgroup variant, interestingly, the form of α leads to two exponentiations modulo p^2 and q^2 with 160-bit exponents. This clearly shows one advantage of this encryption scheme in terms of decryption throughput.

Also, we fix t to 80, that is, we recommend that random numbers r have bitsize 80 or more in practical use.

5 Efficiency

This section gives tight estimates of our cryptosystems' running times for decryption compared to standard ones (OAEP, El Gamal). The elementary unit

will be taken as the number of modular multiplications of bitsize $|n|$ per kilobit of message input; it therefore depends on $|n|$. Typical modulus sizes are $|n| = 512, \cdots, 2048$. We also assume that the execution time of a modular multiplication is quadratic in the operand size and that modular squares are computed by the same routine. Chinese remaindering, as well as random number generation for probabilistic schemes, is considered to be negligible. The parameter t is set to 80 in our two cryptosystems. All secret grandeurs such as factors and exponents are assumed to contain about the same number of ones and zeroes in their binary representation.

We give purely indicative estimates which do not come from actual implementations. Pre-processing stages are not considered, but Chinese remaindering is taken into account whenever possible (hence for all schemes but El Gamal).

Schemes	**Scheme 1**	**Scheme 2**	**OAEP**	**ElGamal**
One-wayness	CR	PDL	RSA	DH
IND-CPA	$D-CR$	$D-PDL$	RSA	D-DH
NM-CCA2	$D-CR$	$D-PDL$	RSA	none
Plaintext size	$\|n\| - 80$	$\|n\| - 80$	$\|n\| - 320$	$\|p\|$
Ciphertext size	$2\ \|n\|$	$2\ \|n\|$	$\|n\|$	$2\ \|p\|$
Decryption Workload (Mult/Kbits)				
$\|n\|, \|p\| = 512$	2731	1707	1024	1536
$\|n\|, \|p\| = 768$	2572	1072	658	1536
$\|n\|, \|p\| = 1024$	2499	781	559	1536
$\|n\|, \|p\| = 1536$	2431	506	485	1536
$\|n\|, \|p\| = 2048$	2398	375	455	1536

6 Conclusion and Further Research

We proposed two new public-key cryptosystems provably semantically secure against adaptive chosen-ciphertext attacks *i.e.* secure in the sense of NM-CCA2. Computationally efficient for decryption, one of them could provide an alternative to OAEP. A typical research topic would be to ensure security against active adversaries relatively to the *computational* related problems CR and PDL. Another (independent) direction consists in improving their decryption throughputs by accelerating computations modulo p^2, possibly using appropriate modular techniques such as [21].

References

1. M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway. Relations Among Notions of Security for Public-Key Encryption Schemes. In *Crypto '98*, LNCS 1462, pages 26–45. Springer-Verlag, 1998.

2. M. Bellare and P. Rogaway. Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In *Proc. of the First ACM CCCS*, pages 62–73. ACM Press, 1993.
3. M. Bellare and P. Rogaway. Optimal Asymmetric Encryption – How to Encrypt with RSA. In *Eurocrypt '94*, LNCS 950, pages 92–111. Springer-Verlag, 1995.
4. D. Bleichenbacher. A Chosen Ciphertext Attack against Protocols based on the RSA Encryption Standard PKCS #1. In *Crypto '98*, LNCS 1462, pages 1–12. Springer-Verlag, 1998.
5. J. S. Coron, D. Naccache and Ju. Stern. A New Signature Forgery Strategy. In *Crypto '99*, Springer-Verlag, 1999.
6. R. Cramer and V. Shoup. A Practical Public Key Cryptosystem Provably Secure against Adaptive Chosen Ciphertext Attack. In *Crypto '98*, LNCS 1462, pages 13–25. Springer-Verlag, 1998.
7. W. Diffie and M. E. Hellman. New Directions in Cryptography. In *IEEE Transactions on Information Theory*, volume IT–22, no. 6, pages 644–654, November 1976.
8. D. Dolev, C. Dwork, and M. Naor. Non-Malleable Cryptography. In *Proc. of the 23rd STOC*. ACM Press, 1991.
9. T. El Gamal. A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms. In *IEEE Transactions on Information Theory*, volume IT–31, no. 4, pages 469–472, July 1985.
10. E. Fujisaki and T. Okamoto. How to Enhance the Security of Public-Key Encryption at Minimum Cost. In *PKC '99*, LNCS 1560, pages 53–68. Springer-Verlag, 1999.
11. S. Goldwasser and S. Micali. Probabilistic Encryption. *Journal of Computer and System Sciences*, 28:270–299, 1984.
12. D. Naccache and J. Stern. A New Cryptosystem based on Higher Residues. In *Proc. of the 5th CCCS*, pages 59–66. ACM press, 1998.
13. M. Naor and M. Yung. Public-Key Cryptosystems Provably Secure against Chosen Ciphertext Attacks. In *Proc. of the 22nd STOC*, pages 427–437. ACM Press, 1990.
14. T. Okamoto and S. Uchiyama. A New Public Key Cryptosystem as Secure as Factoring. In *Eurocrypt '98*, LNCS 1403, pages 308–318. Springer-Verlag, 1998.
15. P. Paillier. Public-Key Cryptosystems Based on Discrete Logarithms Residues. In *Eurocrypt '99*, LNCS 1592, pages 223–238. Springer-Verlag, 1999.
16. D. Pointcheval. New Public Key Cryptosystems based on the Dependent-RSA Problems. In *Eurocrypt '99*, LNCS 1592, pages 239–254. Springer-Verlag, 1999.
17. C. Rackoff and D. R. Simon. Non-Interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack. In *Crypto '91*, LNCS 576, pages 433–444. Springer-Verlag, 1992.
18. R. Rivest, A. Shamir, and L. Adleman. A Method for Obtaining Digital Signatures and Public Key Cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978.
19. RSA Data Security, Inc. Public Key Cryptography Standards – PKCS. Available from `http://www.rsa.com/rsalabs/pubs/PKCS/`.
20. V. Shoup and R. Gennaro. Securing Threshold Cryptosystems against Chosen Ciphertext Attack. In *Eurocrypt '98*, LNCS 1403, pages 1–16. Springer-Verlag, 1998.
21. T. Takagi. Fast RSA-Type Cryptosystems Using N-adic Expansion. In *Crypto '97*, LNCS 1294, pages 372–384. Springer-Verlag, 1997.
22. Y. Tsiounis and M. Yung. On the Security of El Gamal based Encryption. In *PKC '98*, LNCS. Springer-Verlag, 1998.