Proceedings of the 27th International Colloquium on Automata, Languages and Programming (ICALP '2000)
(july 9 - 15, 2000, Geneva, Switzerland) – U. Montanari, J. D. P. Rolim and E. Welzl Eds.
Springer-Verlag, LNCS 1853, pages 499–511.

Extended Notions of Security for Multicast Public Key Cryptosystems

Olivier Baudron, David Pointcheval, and Jacques Stern

École Normale Supérieure, LIENS
45, rue d'Ulm
F-75230 Paris Cedex 05, France
{Olivier.Baudron,David.Pointcheval,Jacques.Stern}@ens.fr

Abstract. In this paper we introduce two notions of security: multi-user indistinguishability and multi-user non-malleability. We believe that they encompass the correct requirements for public key encryption schemes in the context of multicast communications. A precise and non-trivial analysis proves that they are equivalent to the former single-user notions, provided the number of participants is polynomial. We also introduce a new definition for non-malleability which is simpler than those currently in use. We believe that our results are of practical significance: especially they support the use of PKCS#1 v.2 based on OAEP in the multicast setting.
Keywords: Multicast encryption, semantic security, non-malleability.

1 Introduction

1.1 Motivation

With the growth of wide area networks, cryptographic tools often have to coexist and perform related computations. This may raise new security concerns. For example, broadcast encryption has been the subject of several specific attacks, notably directed against low-exponent RSA [20]. Basically, if e is the common public exponent, then e encryptions of a given message under different public keys lead to an easy recovery of the plaintext. Further results by Håstad [14, 22] and Coppersmith [6, 7] proved that "time stamp" variants of broadcast, attaching time to the message before encryption, can be successfully cryptanalyzed with e encrypted messages. So far, most known attacks against RSA assume that related plaintexts have been encrypted to different destinations, which enables an eavesdropper to take advantage of the strong dependences between the RSA permutations, although each one is individually one-way.

Despite these attacks, RSA with small exponents is the de facto standard and multicast encryption is performed in many products by encapsulating a symmetric key within several RSA encryptions together with side data which are specific to each receiver. This is precisely the context that we wish to address and we believe that the related security issues needed to be cleared up in order to ensure confidence in standard designs that allow multicast encryption such as PKCS#1. Thus, albeit technical, our research is of practical significance.

1.2 Notions of security for encryption

In this paper, we wish to propose notions of security that adequately prevent the attacks just mentioned. Usually, a security level is analyzed in terms of the goal and power of an adversary. The ultimate goal that can be achieved is called *invertibility*: given a public key and an encryption of m, retrieve the whole plaintext m. The RSA assumption implies that the basic RSA encryption scheme is non-invertible. As shown in the above example, the related notion dramatically collapses in a broadcast attack. In a different context, stronger notions of security, have been proposed. Goldwasser and Micali define *semantic security* [13] (also called *indistinguishability*) as the inability for an adversary to distinguish encryptions of two plaintexts. This requires probabilistic encryption, where each plaintext has many corresponding ciphertexts, depending on a random parameter. Recent successful attacks against RSA-like cryptosystems [8] based on known plaintext relations stresses the need for proven schemes achieving semantic security.

Surprisingly, the relationship between broadcast attacks and the improved notions of security has not been the subject of specific research, even if known cryptanalyses seem to fail against semantic security. The motivation of this paper is to investigate whether semantic security, contrary to invertibility, is robust in scenarii involving a general notion of multicast. Our first result gives a positive answer: if one can gain a bit of information by considering a specific set of multicast encrypted messages, then at least one scheme used for encryption is not semantically secure. The proof relies on the hybrid technique and is conceptually simple. It is an independant work of Bellare, Boldyreva and Micali who adressed the same problem [1].

Next, we develop a similar analysis with the notion of *non-malleability*, introduced by Dolev, Dwork and Naor [11]. Informally, the notion asserts that, given a ciphertext, it should be impossible to generate a different ciphertext so that the respective plaintexts are related. The problem of encrypted bids is a famous situation where an eavesdropper may try to under-bid a ciphertext of an unknown amount s, without learning anything about s. This is precisely what non-malleability tries to prevent. A broadcast scenario may be envisioned where several recipients collect the bids over a network. The multicast notion requires that the view of many encrypted messages under different public keys gives no advantage in producing the encryption of a related plaintext. Again, we prove that our new definition of multi-user non-malleability is equivalent to the former single-user notion: no broadcast attack can be performed against a non-malleable scheme. Here, the reduction is definitely much harder to obtain. Due to the complex nature of the definitions, involving auxiliary distributions of plaintexts and binary relations, both issued by the attacker, our previous natural reduction cannot be applied. The major technical point of the proof relies on a lemma embedding any distribution into the product of a 2 element-distribution which leads to a simpler definition of non-malleability. We think that this lemma may be of independent interest to cryptographers.

We now discuss the notion of security in terms of the adversary's power. Usually, an attacker is a probabilistic polynomial time Turing machine running in two stages. Firstly, given a public key, it achieves a precomputation stage and halts. From the output data, a challenge is randomly encrypted and given to the attacker which performs a second stage of computation. The polynomial strength of the attacker may be increased by providing him access to a decryption oracle. Whether the oracle is accessible during the first stage only or during whole computation leads to three different scenarii. Under a *chosen-plaintext attack* the adversary can obtain ciphertexts of his choice, which is meaningless in the context of public key encryption. Under *chosen ciphertext attack* [17], the adversary is allowed to use a decryption oracle during the precomputation stage only. Lastly, under *adaptive chosen ciphertext attack* [19], the adversary is allowed to use a decryption oracle during whole algorithm, with the trivial restriction that the challenge cannot be asked to the oracle. The latter is the ideal candidate that one should consider in order to provide the best arguments for security. In our paper, whenever a theorem is stated, it is assumed that one of the three contexts given above has been fixed and hence no decryption oracle is mentioned; potential oracles are preferably viewed as internal parts of the attacker.

1.3 Outline of the paper

The rest of the paper is organized as follows. Section 2 gives common definitions and notations for encryption and probabilities. Sections 3 and 4 contain our analysis of semantic security (which we call indistinguishability) and non-malleability. Both introduce definitions of these notions in the context of multi-cast. The conclusion follows in section 5.

2 Definitions and notations

A public key encryption scheme Π is a triplet $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ consisting of three probabilistic polynomial time algorithms.

- $\mathcal{K}$ is the *key generation algorithm* which, given a security parameter k (usually viewed as a unary input 1^k) produces from its random source ω a pair (pk, sk) of public and secret keys.
- $\mathcal{E}$ is the probabilistic *encryption algorithm* which, given the security parameter k, defines a message space $\mathcal{M}$ such that: for each string x from $\mathcal{M}$, and for each valid public key pk, $\mathcal{E}_{pk}(x)$ is a string y, called the *encryption* of x under pk.
- $\mathcal{D}$ is the (deterministic) *decryption algorithm.* It is required that for every message x in $\mathcal{M}$ and for every pair (pk, sk) output by $\mathcal{K}$, $\mathcal{D}_{sk}(\mathcal{E}_{pk}(x)) = x$. In all other cases, the output of $\mathcal{D}$ is any element of $\mathcal{M} \cup \{\perp\}$. A ciphertext whose decryption is $\perp$ is said to be *invalid.*

A real-valued function $f(n)$ is *negligible* if for any integer k, $|f(n)| < n^{-k}$ for sufficiently large n.

Given a distribution δ over a finite space Ω, we let $\Pr_\delta[E]$ be the probability of an event E. When δ is omitted, it is implicitly assumed that δ is the uniform distribution. The *support* of δ is the set of elements from Ω whose probability is non zero. Often, a random variable is conveniently defined by the output distribution of a probabilistic Turing machine. We let $y \leftarrow TM(x)$ be the result y by running TM on input x and random source ω. If S is a finite set then $y \leftarrow S$ is the operation of picking an element uniformly in S.

When considering several encryption schemes $\Pi_1, .., \Pi_n$ and their related algorithms, we will denote by $\mathcal{K}^n$, $\mathcal{E}^n$ and $\mathcal{D}^n$ the algorithms that given an input vector of n adequate data, output a vector of dimension n whose distribution is given by the product of the output distributions of $\mathcal{K}_1 \times .. \times \mathcal{K}_n$, $\mathcal{E}_1 \times .. \times \mathcal{E}_n$ and $\mathcal{D}_1 \times .. \times \mathcal{D}_n$ respectively. We insist that all encryption schemes need not be identical.

Our multicast notion enlarges the intuitive definition of broadcast when a unique plaintext is encrypted. In this paper, we consider a multicast communication as a set of encryptions of suitably related plaintexts under different public keys. For example the reader might consider messages containing the name of the recipient followed by a possibly common text. Formally, a broadcast distribution of plaintexts is any *diagonal* distribution whose support is in $\mathcal{M}^n$ whereas a multicast distribution of plaintexts is any distribution whose support is in $\mathcal{M}^n$.

3 Indistinguishability

3.1 Single-user encryption schemes

Secure encryption should preserve privacy even in the critical context where the messages are taken from a small set of plaintexts: it should be impossible for an eavesdropper to distinguish encryptions of distinct values. Such a requirement is captured by the notion of indistinguishability, also known as semantic security [13, 15]. Examples, secure against chosen plaintext attack, include El Gamal [12] (based on the decisional Diffie-Hellman assumption [10]), Naccache-Stern [16] (based on higher residues) and Okamoto-Uchiyama [18] (based on factorization). Our definition exactly follows [2] and uses the same notations. Indistinguishability is defined by the advantage of an adversary $A = (A_1, A_2)$ performing a sequence of two algorithms.

In a first step, algorithm A_1 is run on input of the public key pk and outputs two plaintexts messages x^0 and x^1 plus a string s encoding information to be handled to A_2. Next a message from $\{x^0, x^1\}$ is chosen at random and encrypted into a challenge ciphertext y. In a second step, A_2 is given the input (y, s) and has to guess the bit of the plaintext being encrypted. The advantage of A is measured by the probability that it outputs the correct bit of the challenge. The scheme is indistinguishable if no adversary obtains an advantage significantly greater than one would obtain by flipping a coin. The formal definition follows:

Definition 1. *Single-user indistinguishability.*
Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme with a security parameter k and let $A = (A_1, A_2)$ be an adversary. For $k \in \mathbb{N}$, we define the advantage:

$$\begin{aligned}\mathsf{Adv}_{A,\Pi}(k) = 2 \Pr \big[&(pk, sk) \leftarrow \mathcal{K}(1^k);\ (x^0, x^1, s) \leftarrow A_1(pk);\ b \leftarrow \{0,1\};\\ &y \leftarrow \mathcal{E}_{pk}(x^b)\ :\ A_2(s, y) = b \big] - 1\end{aligned}$$

We say that Π is single-user indistinguishable *(S-IND)* if for every polynomial time adversary A, $\mathsf{Adv}_{A,\Pi}(k)$ is negligible.

3.2 Multicast encryption schemes

In the context of multicast, the usual notion of indistinguishability does not, by itself, guarantee that no bit of information is leaked when putting together the encryptions of related messages under different public keys. Our definition captures this stronger notion of security by giving the adversary the ability to choose two vectors of plaintexts whose coordinates are plaintext messages possibly related or even identical. Next, one of the two vectors is chosen at random and is encrypted coordinatewise with the different public keys. The final goal of the adversary is to guess which one was encrypted. This is easily done if a boolean function distinguishes the two vectors of plaintexts and is computable from the encrypted data. Again our formal definition is in terms of the advantage of an adversary playing the game just given. In the following, underlined variables denote vectors of size n; the i^{th} coordinate refers to the i^{th} cryptosystem.

Definition 2. *Multi-user indistinguishability.*
Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme with a security parameter k and let $A = (A_1, A_2)$ be an adversary. For $k, n \in \mathbb{N}$, we define the advantage:

$$\mathsf{Adv}_{A,\Pi}(k,n) = 2\Pr\Big[(\underline{pk}, \underline{sk}) \leftarrow \mathcal{K}^n(1^k);\ (\underline{x}^0, \underline{x}^1, s) \leftarrow A_1(\underline{pk});\ b \leftarrow \{0,1\};\\ \underline{y} \leftarrow \mathcal{E}_{\underline{pk}}(\underline{x}^b)\ :\ A_2(s, \underline{y}) = b\Big] - 1$$

We say that Π is multi-user indistinguishable *(M-IND)* if for every polynomial time adversary A, $\mathsf{Adv}_{A,\Pi}(k,n)$ is negligible.

3.3 Results

As expected, any multi-user indistinguishable encryption scheme Π is also single-user indistinguishable. Indeed, if an adversary distinguishes $\mathcal{E}_{pk}(m^0)$ from $\mathcal{E}_{pk}$ (m^1) then it obviously distinguishes two encrypted vectors whose first coordinate is the encryption of m^0 and m^1 under the public key pk. Also note that the usual definition of (single-user) indistinguishability, expressed in [2], is the particular case of multi-user indistinguishability where $n = 1$. The following result achieves equivalence.

Theorem 3. *S-IND⇒M-IND.*
If encryption scheme Π is single-user indistinguishable, then it is multi-user indistinguishable.

Proof. Let A be an adversary attacking Π in the sense of M-IND. We build n adversaries $B_i = (B_{i,1}, B_{i,2})_{1 \leq i \leq n}$, as follows:

Algorithm $B_{i,1}(pk_i)$:
  $\underline{pk} \leftarrow (pk_1, .., pk_i, .., pk_n)$
  $(\underline{x}^0, \underline{x}^1, s) \leftarrow A_1(\underline{pk})$
  return(x_i^0, x_i^1, s)

Algorithm $B_{i,2}(y_i, s)$:

$b' \leftarrow \{0,1\}$
$\underline{y} \leftarrow (y_1, .., y_i, .., y_n)$ with $y_j = \mathcal{E}_{pk_j}(x_j^{b'})$ if $j < i$
$\qquad\qquad\qquad\qquad\qquad\quad y_j = \mathcal{E}_{pk_j}(x_j^{\bar{b}'})$ if $j > i$
$b'' \leftarrow A_2(\underline{y}, s)$
$\textsf{return}\ b''$

In a first step $B_{i,1}$ extends pk_i to a vector of public keys $\underline{pk}$, using $(n-1)$ times the algorithm $\mathcal{K}$. Then A_1 is run with the input $\underline{pk}$. The i^{th} pair of plaintext messages output by A_1 is returned, which completes the first part of the algorithm. We note b the unknown bit of the challenge, i.e. $y_i = \mathcal{E}_{pk_i}(x_i^b)$. In a second step, $B_{i,2}$ extends its input y_i to a hybrid vector $\underline{y}$: the first coordinates of $\underline{y}$ come from the encryption of $\underline{x}^{b'}$ whereas the last coordinates of $\underline{y}$ come from the encryption of $\underline{x}^{\bar{b}'}$. Bit b'' output by running A_2 on $\underline{y}$ is returned as an answer to the challenge.

We now compute the advantage of B_i for $\underline{pk}$, $\underline{x}^0$, $\underline{x}^1$ and s fixed. Let d be a random bit and let $\Pr_i$ (respectively $\Pr'_i$) be the probability that the initial adversary A_2 successfully guesses the plaintext of the left (respectively right) part of a hybrid ciphertext formed with i coordinates from x^d followed by $(n-i)$ coordinates from $x^{\bar{d}}$:

$$\Pr_i = \Pr\left[d \leftarrow \{0,1\};\ \underline{c} \leftarrow \mathcal{E}_{\underline{pk}}(a_1^d, .., a_i^d, a_{i+1}^{\bar{d}}, .., a_n^{\bar{d}});\ d' \leftarrow A_2(\underline{c}, s)\ :\ d' = d\right]$$
$$\Pr'_i = \Pr\left[d \leftarrow \{0,1\};\ \underline{c} \leftarrow \mathcal{E}_{\underline{pk}}(a_1^d, .., a_i^d, a_{i+1}^{\bar{d}}, .., a_n^{\bar{d}});\ d' \leftarrow A_2(\underline{c}, s)\ :\ d' \neq d\right]$$

Note that,

$$\Pr_i + \Pr'_i = 1 \tag{1}$$

We apply Bayes' theorem, considering the value of the bit b' randomly chosen in the algorithm $B_{i,2}$:

$$\begin{aligned}
&\Pr\left[b \leftarrow \{0,1\};\ y_i \leftarrow \mathcal{E}_{pk_i}(x_i^b);\ b'' \leftarrow B_{i,2}(y_i, s)\ :\ b'' = b\right] \\
&= \tfrac{1}{2}\Pr\left[b \leftarrow \{0,1\};\ y_i \leftarrow \mathcal{E}_{pk_i}(x_i^b);\ b'' \leftarrow B_{i,2}(y_i, s)\ :\ b'' = b \mid b' = b\right] \\
&\quad + \tfrac{1}{2}\Pr\left[b \leftarrow \{0,1\};\ y_i \leftarrow \mathcal{E}_{pk_i}(x_i^b);\ b'' \leftarrow B_{i,2}(y_i, s)\ :\ b'' = b \mid b' \neq b\right] \\
&= \tfrac{1}{2}\Pr_i + \tfrac{1}{2}\Pr'_{i-1} \qquad (2)
\end{aligned}$$

It follows from (1) and (2) that the advantage of B_i is:

$$\mathsf{Adv}_{B_i,\Pi} = 2\left(\tfrac{1}{2}\Pr_i + \tfrac{1}{2}\Pr'_{i-1}\right) - 1 = \Pr_i - \Pr_{i-1}$$

Middle terms cancel in the sum, so that:

$$\sum_{i=1}^{n} \mathsf{Adv}_{B_i,\Pi} = \Pr_n - \Pr_0 = \mathsf{Adv}_{A,\Pi}$$

Consequently, if i is uniformly chosen at random in $\{1, .., n\}$, we obtain a reduction from a multi-distinguisher attacker A with advantage ϵ, to a single-distinguisher attacker B with advantage ϵ/n. □

4 Non-malleability

4.1 Single-user non-malleability

The notion of non-malleability was introduced in [11] and formalized in a different manner in [2]. The main idea is that, given an encrypted message y, an adversary is unable to output a ciphertext y' whose decryption is related to the decryption of y. More precisely, this goes along an interactive experiment with an adversary $A = (A_1, A_2)$ which is described below.

The Turing machine A_1 is run with input of a public key pk and outputs the description of a probabilistic polynomial time Turing machine M, and a string s for further computation. The output of M defines a distribution of plaintext messages whose support is a set $|M| \subset \mathcal{M}$. In the following M refers to the Turing machine as well as its output distribution. Then a message x is randomly chosen by running M and its encryption is given to A_2. The goal of A_2 is to output a binary relation R over $|M| \times \mathcal{M}$ and a ciphertext $y' \neq y$ whose decryption x' is related to x according to R. The scheme is non-malleable if for any adversary the probability that $R(x, x')$ holds is not significantly better than the probability that $R(\tilde{x}, x')$ for a random $\tilde{x}$ from M.

For notational convenience we have simplified the definition given in [2]. In the original paper, the goal of the adversary was to output a vector $\mathbf{y}'$ of $t-1$ ciphertexts related to y according to a relation R of arity t. In this case, it is required that no coordinate of $\mathbf{y}'$ is equal to y. It was also proven that both definitions were not equivalent. The former could not be reduced to the latter. In the rest of our paper we will only represent elements y' with one coordinate so that no confusion arises with vectors from the broadcast notation. But one can also build a similar theory of multi-user non-malleability for relations of arity t by considering the modified ciphertext as a vector of ciphertext vectors $\underline{\mathbf{y}}'$ and an appropriate binary relation over $|M| \times \mathcal{M}^{n\times(t-1)}$.

Recently, it was shown by Bellare and Sahai [4] that non-malleability (in any attack model) was equivalent to indistinguishability where the adversary gets the additional power of "parallel ciphertext attack" (i.e. non adaptive ciphertext attack after seeing the challenge encryption). Consequently, our first result may apply to this notion. However, we followed the standard definition of non-malleability and proved it may be simplified.

4.2 Multi-user non-malleability

Scenarii where it is unclear whether single-non-malleability is enough to ensure a satisfactory notion of security can be envisioned: for example, the view of different encryptions under several public keys might give the opportunity for an adversary to flip one of the encrypted message into its opposite. It is also not clear that encrypted messages sent to different users may not be exchanged. Thus, if one wishes to cover the standard context of multicast it is natural to give an extended notion of security for non-malleability which we now undertake.

The adversary is given n public keys and outputs a probabilistic polynomial time Turing machine M plus a string s. By running M on a random source we require that its output defines a distribution of plaintext messages whose support $|M|$ is in $\mathcal{M}^n$. Then, a vector $\underline{x}$ is randomly chosen by running M, and its coordinatewise encryption according to the different public keys is given to A_2. The goal of A_2 is to output a vector of ciphertexts $\underline{y}'$ and a relation R over $|M| \times \mathcal{M}^n$. A is successful if R relates the corresponding decrypted messages. The formal definition is given below.

Remark. The exact support $|M|$ of M may not be computable in polynomial time. It is therefore only required that the relation R is defined on a subset of $\mathcal{M}^n \times \mathcal{M}^n$ and covers $|M| \times \mathcal{M}^n$.

Definition 4. *Multi-user non-malleability.*
Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme with security parameter k and let $A = (A_1, A_2)$ be an adversary. For $k, n \in \mathbb{N}$, we define the advantage:

$$\mathsf{Adv}_{A,\Pi}(k,n) = \left|\mathsf{Succ}_{A,\Pi}(k,n) - \mathsf{Succ}_{A,\Pi,\$}(k,n)\right|,$$

where

$$\begin{aligned}\mathsf{Succ}_{A,\Pi}(k,n) =& \Pr\big[\,(\underline{pk},\underline{sk}) \leftarrow \mathcal{K}^n(1^k);\ (M,s) \leftarrow A_1(\underline{pk});\ \underline{x} \leftarrow M;\ \underline{y} \leftarrow \mathcal{E}_{\underline{pk}}(\underline{x});\\ & (R,\underline{y}') \leftarrow A_2(M,s,\underline{y});\ \underline{x}' \leftarrow \mathcal{D}_{\underline{sk}}(\underline{y}')\ :\ \perp \notin \underline{x}' \wedge R(\underline{x},\underline{x}')\big]\end{aligned}$$

$$\begin{aligned}\mathsf{Succ}_{A,\Pi,\$}(k,n) =& \Pr\big[\,(\underline{pk},\underline{sk}) \leftarrow \mathcal{K}^n(1^k);\ (M,s) \leftarrow A_1(\underline{pk});\ \underline{x},\underline{\tilde{x}} \leftarrow M;\ \underline{y} \leftarrow \mathcal{E}_{\underline{pk}}(\underline{x});\\ & (R,\underline{y}') \leftarrow A_2(M,s,\underline{y});\ \underline{x}' \leftarrow \mathcal{D}_{\underline{sk}}(\underline{y}')\ :\ \perp \notin \underline{x}' \wedge R(\underline{\tilde{x}}',\underline{x}')\big]\end{aligned}$$

with $\tilde{x}'_i = \begin{cases} x_i \text{ if } & y'_i = y_i \\ \tilde{x}_i \text{ if } & y'_i \neq y_i \end{cases}$, for each i in $\{1,..,n\}$

We say that Π is multi-user non-malleable (M-NM) if for every polynomial time adversary A whose output is a distribution of plaintexts M and a relation R both computable in polynomial time then $\mathsf{Adv}_{A,\Pi}$ is negligible.

The motivation to introduce a new variable $\underline{\tilde{x}}'$ was to restrict the domain of the random variable $\underline{\tilde{x}}$ for the coordinates left unchanged by A_2. This is the condition in dimension n of the requirement $y' \neq y$ in dimension 1, defined in [2]. This rule makes the adversary gain no advantage in partially copying a vector of ciphertexts and outputting a relation whose value is true on domains of the form $((x_0,..,*),(x_0,..,*))$.

The usual notion of (single-user) non-malleability is the particular case where n is fixed to 1.

4.3 Results

The next result is the main technical achievement of our paper and leads to a simplified definition of non-malleability. It claims that the distribution of plaintexts M can be restricted to an atomic form.

Lemma 5. *Atomic non-malleability.*
Let Π be an encryption scheme and let A be an adversary attacking Π in the

sense of M-NM. Then there exists another adversary B attacking Π, in the sense of M-NM such that the distribution of plaintexts that B outputs is always a uniform distribution of two vectors of plaintexts. Moreover, the running time of B is that of A plus the running time of the Turing machine M output by A.

Proof. The adversary $B = (B_1, B_2)$ is defined as follows:

Algorithm $B_1(\underline{pk})$	Algorithm $B_2(\underline{y}, s)$
$(M, s) \leftarrow A_1(\underline{pk})$	$(R, \underline{y}') \leftarrow A_2(\underline{y}, s)$
$\underline{a}^0 \leftarrow M;\ \underline{a}^1 \leftarrow M$	return $(R, \underline{y}')$
return $(\{\underline{a}^0, \underline{a}^1\}, s)$	

Here the description of B_2 is identical to A_2 except that the relation R is restricted to the set $\{\underline{a}^0, \underline{a}^1\} \times \mathcal{M}^n$ instead of $M \times \mathcal{M}^n$. We first claim that the input distribution of the ciphertexts is the same for A_2 and B_2. Indeed, using Bayes' theorem and since $\underline{x}$ has equal probability $1/2$ of being $\underline{a}_0$ or $\underline{a}_1$, it results that for all $\underline{X}$ in M:

$$\begin{aligned}
&\Pr[\underline{a}^0, \underline{a}^1 \leftarrow M;\ \underline{x} \leftarrow \{\underline{a}^0, \underline{a}^1\}\ :\ \underline{x} = \underline{X}] \\
&\quad = \tfrac{1}{2}\Pr\left[\underline{a}^0 \leftarrow M\ :\ \underline{a}^0 = \underline{X}\right] + \tfrac{1}{2}\Pr\left[\underline{a}^1 \leftarrow M\ :\ \underline{a}^1 = \underline{X}\right] \\
&\quad = \Pr[\underline{x} \leftarrow M\ :\ \underline{x} = \underline{X}]
\end{aligned}$$

Consequently, $\mathsf{Succ}_{B,\Pi} = \mathsf{Succ}_{A,\Pi}$. Next, in order to express $\mathsf{Succ}_{B,\Pi,\$}$ we decorelate $\underline{\tilde{x}}$ from $\underline{x}$, considering its two possible values among $\{\underline{a}^0, \underline{a}^1\}$. Using the notations from definition 3, it holds:

$$\begin{aligned}
&\Pr\left[\underline{a}^0, \underline{a}^1 \leftarrow M;\ \underline{x}, \underline{\tilde{x}} \leftarrow \{\underline{a}^0, \underline{a}^1\}\ :\ R(\underline{\tilde{x}}', \underline{x}')\right] \\
&= \tfrac{1}{2}\Pr[\underline{a}^0, \underline{a}^1 \leftarrow M;\ \underline{x}, \underline{\tilde{x}} \leftarrow \{\underline{a}^0, \underline{a}^1\}\ :\ R(\underline{\tilde{x}}', \underline{x}') \mid \underline{\tilde{x}} = \underline{x}] \\
&\quad + \tfrac{1}{2}\Pr[\underline{a}^0, \underline{a}^1 \leftarrow M;\ \underline{x}, \underline{\tilde{x}} \leftarrow \{\underline{a}^0, \underline{a}^1\}\ :\ R(\underline{\tilde{x}}', \underline{x}') \mid \underline{\tilde{x}} \neq \underline{x}] \\
&= \tfrac{1}{2}\Pr[\underline{a}^0, \underline{a}^1 \leftarrow M;\ \underline{x} \leftarrow \{\underline{a}^0, \underline{a}^1\}\ :\ R(\underline{x}, \underline{x}')] \\
&\quad + \tfrac{1}{2}\Pr\left[\underline{\tilde{a}}^0, \underline{a}^1 \leftarrow M\ :\ R(\underline{\tilde{a}}^{0\prime}, \underline{a}^{1\prime})\right]
\end{aligned}$$

So, $\mathsf{Succ}_{B,\Pi,\$} = \frac{1}{2}\mathsf{Succ}_{B,\Pi} + \frac{1}{2}\mathsf{Succ}_{A,\Pi,\$}$ and $\mathsf{Adv}_{B,\Pi} = \mathsf{Succ}_{B,\Pi} - \mathsf{Succ}_{B,\Pi,\$} = \frac{1}{2}\mathsf{Succ}_{B,\Pi} - \frac{1}{2}\mathsf{Succ}_{A,\Pi,\$}$. With the previous result, we conclude

$$\mathsf{Adv}_{B,\Pi} = \frac{1}{2}\mathsf{Adv}_{A,\Pi}$$

□

It is easily seen that the definition of single-user non-malleability is the restricted case of the multi-user non-malleability for $n = 1$. The equivalence follows from the next result.

Theorem 6. *S-NM⇒M-NM. If encryption scheme Π is single-user non-malleable, then it is multi-user non-malleable.*

Proof. Let $A = (A_1, A_2)$ be an adversary attacking Π in the sense of multi-user non-malleability with an advantage ϵ. Without loss of generality, as was shown in Lemma 1, we assume that A_1 outputs a uniform distribution M of two plaintext vectors $\underline{a}_0$ and $\underline{a}_1$. We will build n Turing machine $B_1, .., B_n$ attacking Π in the sense of single-user non-malleability. For any $i \in \{1, .., n\}$, the description of $B_i = (B_{i,1}, B_{i,2})$ is as follows:

Algorithm $B_{i,1}(pk_i)$:
  $\underline{pk} \leftarrow (pk_1, .., pk_i, .., pk_n)$
  $(M, s) \leftarrow A_1(\underline{pk})$
  return $M_i = \{a_i^0, a_i^1\}$

Algorithm $B_{i,2}(c_i, s)$:
  $b' \leftarrow \{0, 1\}$
  $\underline{c} \leftarrow (c_1, .., c_i, .., c_n)$ with $c_j = \mathcal{E}_{pk_j}(a_j^{b'})$ if $j < i$
  $c_j = \mathcal{E}_{pk_j}(a_j^{\bar{b}'})$ if $j > i$
  $(\underline{c}', R) \leftarrow A_2(\underline{c}, s)$
  $R_i(a_i^k, u) \iff R(\underline{a}^k, \underline{v})$ with $v_i = u$
  $v_j = \mathcal{D}_{sk_j}(c'_j)$ if $j \neq i$
  return $({c'}_i, R_i)$

As in the previous construction, the first part of the algorithm extends the input pk_i into a vector $\underline{pk}$ and calls the attacker A_1 on this data. Without loss of generality, as was shown in Lemma 1, A_1 outputs a distribution M of two plaintexts $\underline{a}_0$ and $\underline{a}_1$. Then both i^{th} coordinates are returned. The algorithm $B_{i,2}$ takes as input the ciphertext c_i of a plaintext a_i^b where b is an unknown bit. We focus on the way the binary relation R_i over $\{a_i^0, a_i^1\} \times \mathcal{M}$ is built from the initial relation R over $\{\underline{a}_0, \underline{a}_1\} \times \mathcal{M}^n$. Since the expression of the advantage of A only depends on the decryption of $\underline{c}$, we let the i^{th} coordinate free and fix the others to the decrypted coordinates of $\underline{c}$ thanks to the knowledge of the related secret keys. Thus R_i is the section of R on this particular sub-space. Note that, the exact definition of R_i may be ambiguous in the case where $a_i^0 = a_i^1$ and $\underline{a}^0 \neq \underline{a}^1$. Here, it is clear that any attacker (even infinitely powerful) obtains a null advantage since the encryption of a_i^b is perfectly independent of the bit b. Thus in this specific case, the definition of R_i has little importance, and for convenience, it is defined by choosing b randomly so that the following computations remain true.

We now fix $\underline{pk}$, $\underline{a}_0$ and $\underline{a}_1$. The main goal is to analyze the behavior of the adversary A_2 when its input is a hybrid vector of ciphertexts from $\underline{a}_0$ and $\underline{a}_1$. Let Pr_i (respectively ${\mathrm{Pr}'}_i$) be the probability that A_2 successfully outputs a ciphertext related to the first (respectively last) part of the initial hybrid plaintext.

$$\mathrm{Pr}_i = \mathrm{Pr}\left[b \leftarrow \{0,1\};\ \underline{c} \leftarrow \mathcal{E}_{\underline{pk}}(a_1^b,..,a_i^b,a_{i+1}^{\bar{b}},..,a_n^{\bar{b}});\ (\underline{c}',R) \leftarrow A_2(\underline{c},s) : R(\underline{a}^b,\mathcal{D}_{\underline{sk}}(\underline{c}'))\right]$$
$$\mathrm{Pr}'_i = \mathrm{Pr}\left[b \leftarrow \{0,1\};\ \underline{c} \leftarrow \mathcal{E}_{\underline{pk}}(a_1^b,..,a_i^b,a_{i+1}^{\bar{b}},..,a_n^{\bar{b}});\ (\underline{c}',R) \leftarrow A_2(\underline{c},s) : R(\underline{a}^{\bar{b}},\mathcal{D}_{\underline{sk}}(\underline{c}'))\right]$$

Remark: If $a_i^0 = a_i^1$ then a_i^b can be linked identically to the left part or the right part of the hybrid, hence $\mathrm{Pr}_i = \mathrm{Pr}_{i-1}$ and $\mathrm{Pr}'_i = \mathrm{Pr}'_{i-1}$.

It follows from the above definitions that $\Pr_n = \Pr'_0$ and $\Pr'_n = \Pr_0$.
The success of the attacker B_i is:

$$
\begin{aligned}
&\mathsf{Succ}_{B_i,\Pi} \\
&= \Pr\left[b,b' \leftarrow \{0,1\};\ \underline{c} \leftarrow (c_1,..,c_i,..,c_n);\ (\underline{c}',R) \leftarrow A_2(\underline{c},s)\ :\ R(\underline{a}^b,\mathcal{D}_{\underline{pk}}(\underline{c}'))\right] \\
&= \tfrac{1}{2}\Pr\left[b,b' \leftarrow \{0,1\};\ \underline{c} \leftarrow (c_1,..,c_n);\ (\underline{c}',R) \leftarrow A_2(\underline{c},s)\ :\ R(\underline{a}^b,\mathcal{D}_{\underline{pk}}(\underline{c}')) \mid b' = b\right] \\
&\quad + \tfrac{1}{2}\Pr\left[b,b' \leftarrow \{0,1\};\ \underline{c} \leftarrow (c_1,..,c_n);\ (\underline{c}',R) \leftarrow A_2(\underline{c},s)\ :\ R(\underline{a}^b,\mathcal{D}_{\underline{pk}}(\underline{c}')) \mid b' \neq b\right] \\
&= \tfrac{1}{2}\Pr_i\ +\ \tfrac{1}{2}\Pr'_{i-1}
\end{aligned}
$$

The average success Succ is obtained by considering the four possible values of the B-bit b' and the random bit $\tilde{b}$ relatively to the challenge bit b. Since b shares the vector $\underline{c}$ into a left part of $i-1$ encrypted coordinates from b' and a right part of $(n-1-i)$ encrypted coordinates from $\bar{b}'$, whether b is equal to b' or $\bar{b}'$ leads to an hybrid vector $\underline{c}$ whose frontier is at position i or $i-1$. In each case, whether the random bit $\tilde{b}$ is the left or the right part of the hybrid vector $\underline{c}$, leads to one of the expressions $\Pr$ or $\Pr'$.

Let the distribution: $\delta = \left\{b,b',\tilde{b} \leftarrow \{0,1\};\ \underline{c} \leftarrow (c_1,..,c_i,..,c_n);\ (\underline{c}',R) \leftarrow A_2(\underline{c},s)\right\}$.

$$
\begin{aligned}
&\mathsf{Succ}_{B_i,\Pi,\$} \\
&= \Pr_\delta\left[R(\underline{a}^{\tilde{b}},\mathcal{D}_{\underline{pk}}(\underline{c}'))\right] \\
&= \tfrac{1}{4}\Pr_\delta\left[R(\underline{a}^{\tilde{b}},\mathcal{D}_{\underline{pk}}(\underline{c}')) \mid \tilde{b}=b \wedge b'=b\right]\ +\ \tfrac{1}{4}\Pr_\delta\left[R(\underline{a}^{\tilde{b}},\mathcal{D}_{\underline{pk}}(\underline{c}')) \mid \tilde{b}=b \wedge b'\neq b\right] \\
&\quad + \tfrac{1}{4}\Pr_\delta\left[R(\underline{a}^{\tilde{b}},\mathcal{D}_{\underline{pk}}(\underline{c}')) \mid \tilde{b}\neq b \wedge b'=b\right]\ +\ \tfrac{1}{4}\Pr_\delta\left[R(\underline{a}^{\tilde{b}},\mathcal{D}_{\underline{pk}}(\underline{c}')) \mid \tilde{b}\neq b \wedge b'\neq b\right] \\
&= \tfrac{1}{4}\Pr_i\ +\ \tfrac{1}{4}\Pr'_{i-1} + \tfrac{1}{4}\Pr'_i + \tfrac{1}{4}\Pr'_{i-1}
\end{aligned}
$$

It follows that the advantage of B_i is:

$$\mathsf{Adv}_{B_i} = \mathsf{Succ}_{B_i,\Pi} - \mathsf{Succ}_{B_i,\Pi,\$} = \tfrac{1}{4}\Pr_i + \tfrac{1}{4}\Pr'_{i-1} - \tfrac{1}{4}\Pr'_i - \tfrac{1}{4}\Pr_{i-1}$$

Remark: if $a_i^0 = a_i^1$ then from the previous remark $Adv_{B_i} = 0$ as expected.

Finally the sum is:

$$\sum_{i=1}^{n} \mathsf{Adv}_{B_i} = \tfrac{1}{4}(\Pr_n + \Pr'_0 - \Pr'_n - \Pr_0) = \tfrac{1}{2}(\Pr_n - \Pr_0) = \mathsf{Adv}_A$$

Thus, if i is randomly choosen in the set $\{1,..,n\}$, one obtains a reduction from a global adversary with advantage ϵ to an adversary with advantage ϵ/n against a single cryptosystem. $\square$

Consequences of the results. In the case of adaptive chosen ciphertext attacks, it was proved by Bellare *et al.* [2] that both notions of indistinguishability and non-malleability are equivalent, and hence are also equivalent to the multi-user notions of security. Thus, our results show that some recent encryption schemes achieve a high level of multicast security requirement. In the random oracle model, one can mention the RSA-base OAEP [3] from Bellare and Rogaway. It was recently adopted as a standard of encryption in the PKCS#1 [21, 5] specifications. In the standard model of proofs, only the Cramer-Shoup scheme [9] achieves proven security and practical effectiveness. Finally, we point out some practical and straightforward applications of multi-user secure encryption. This includes pay-per-view television, where a part of the bandwith is used to broadcast encrypted keys to each user. Secure electronic mail such as PGP is also given better confidence especially when adressing several recipients. One may also envision secure election protocols with a large number of independent authorities generally resulting in many related encrypted plaintexts. Lastly, multi-party computations usually use the assumption of a broadcast channel and thus should benefit from our multicast notions of secutity.

5 Conclusion

We have extended the applicability of two powerful notions of security: indistinguishability and non-malleability. Every known attack is now covered by our new multicast security definitions. Furthermore, the reductions that we have shown have linear coefficients in the number of users. As a consequence, we believe that proven encryptions schemes with common single-user security parameters are ready to be safely spread over the Internet.

Acknoledgments

We thanks the program commitee for their valuable comments.

References

1. M. Bellare, A. Boldyreva, and S. Micali. Public-Key Encryption in a Multi-user Setting : Security Proofs and Improvements. In *Eurocrypt '00*, LNCS. Springer-Verlag, 2000.
2. 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.
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. M. Bellare and A. Sahai. Non-Malleable Encryption : Equivalence between Two Notions and an Indistinguishability-Based Characterization. In *Crypto '99*, LNCS 1666, pages 519–536. Springer-Verlag, 1998.
5. 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.
6. D. Coppersmith. Finding a Small Root of a Univariate Modular Equation. In *Eurocrypt '96*, LNCS 1070, pages 155–165. Springer-Verlag, 1996.
7. D. Coppersmith. Small Solutions to Polynomial Equations, and Low Exponent RSA Vulnerabilities. *Journal of Cryptology*, 10:233–260, 1997.
8. D. Coppersmith, M. Franklin, J. Patarin, and M. Reiter. Low-Exponent RSA with Related Messages. In *Eurocrypt '96*, LNCS 1070, pages 1–9. Springer-Verlag, 1996.
9. 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.
10. 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.
11. D. Dolev, C. Dwork, and M. Naor. Non-Malleable Cryptography. In *Proc. of the 23rd STOC*. ACM Press, 1991.

12. 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.
13. S. Goldwasser and S. Micali. Probabilistic Encryption. *Journal of Computer and System Sciences*, 28:270–299, 1984.
14. J. Håstad. Solving Simultaneous Modular Equations of Low Degree. *SIAM Journal of Computing*, 17:336–341, 1988.
15. S. Micali, C. Rackoff, and R. Sloan. The notion of security for probabilistic cryptosystems. *SIAM J. of Computing*, April 1988.
16. D. Naccache and J. Stern. A New Cryptosystem based on Higher Residues. In *Proc. of the 5th CCCS*, pages 59–66. ACM press, 1998.
17. 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.
18. 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.
19. 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.
20. 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.
21. RSA Data Security, Inc. Public Key Cryptography Standards – PKCS. Available from `http://www.rsa.com/rsalabs/pubs/PKCS/`.
22. H. Shimizu. On the Improvement of the Håstad Bound. In *1996 IEICE Fall Conference*, Volume A-162, 1996. In Japanese.