

Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm

Mihir Bellare and Chanathip Namprempre

Dept. of Computer Science & Engineering, University of California at San Diego,
9500 Gilman Drive, La Jolla, California 92093, USA.
{mihir, cnamprem}@cs.ucsd.edu
www-cse.ucsd.edu/users/{mihir, cnamprem}

Abstract. We consider two possible notions of authenticity for symmetric encryption schemes, namely integrity of plaintexts and integrity of ciphertexts, and relate them to the standard notions of privacy for symmetric encryption schemes by presenting implications and separations between all notions considered. We then analyze the security of authenticated encryption schemes designed by “generic composition,” meaning making black-box use of a given symmetric encryption scheme and a given MAC. Three composition methods are considered, namely *Encrypt-and-MAC plaintext*, *MAC-then-encrypt*, and *Encrypt-then-MAC*. For each of these, and for each notion of security, we indicate whether or not the resulting scheme meets the notion in question assuming the given symmetric encryption scheme is secure against chosen-plaintext attack and the given MAC is unforgeable under chosen-message attack. We provide proofs for the cases where the answer is “yes” and counter-examples for the cases where the answer is “no.”

1 Introduction

We use the term *authenticated encryption scheme* to refer to a shared-key based transform whose goal is to provide *both* privacy *and* authenticity of the encapsulated data. In such a scheme the *encryption* process applied by the sender takes the key and a plaintext to return a ciphertext, while the *decryption* process applied by the receiver takes the same key and a ciphertext to return either a plaintext or a special symbol indicating that it considers the ciphertext invalid or unauthentic.

The design of such schemes has attracted a lot of attention historically. The early schemes were typically based on adding “redundancy” to the message before CBC encrypting, and many of these schemes were broken. Today authenticated encryption schemes continue to be the target of design and standardization efforts. A popular modern design paradigm is to combine MACs with standard block cipher modes of operation.

The goal of symmetric encryption is usually viewed as privacy, but an authenticated encryption scheme is simply a symmetric encryption scheme meeting additional authenticity goals. The first part of this paper formalizes several different possible notions of authenticity for symmetric encryption schemes, and integrates them into the existing mosaic of notions by relating them to the main known notions of privacy for

symmetric encryption, via implications and separations in the style of [3]. The second part of this paper is motivated by emerging standards such as [16] which design authenticated encryption schemes by what we call “generic composition” of encryption and MAC schemes. We analyze, with regard to meeting the previous notions, several generic composition methods. Let us now look at these items in more detail.

1.1 Relations among Notions

Privacy goals for symmetric encryption schemes include indistinguishability and non-malleability, each of which can be considered under either chosen-plaintext or (adaptive) chosen-ciphertext attack, leading to four notions of security we abbreviate IND-CPA, IND-CCA, NM-CPA, NM-CCA. (The original definitions were in the asymmetric setting [12,10,18] but can be “lifted” to the symmetric setting using the encryption oracle based template of [2]). The relations among these notions are well-understood [3,11]. (These papers state results for the asymmetric setting, but as noted in [3] it is an easy exercise to transfer them to the symmetric setting.)

We consider two notions of integrity (we use the terms authenticity and integrity interchangeably) for symmetric encryption schemes. INT-PTXT (integrity of plaintexts) requires that it be computationally infeasible to produce a ciphertext decrypting to a message which the sender had never encrypted, while INT-CTXT (integrity of ciphertexts) requires that it be computationally infeasible to produce a ciphertext not previously produced by the sender, regardless of whether or not the underlying plaintext is “new.” (In both cases, the adversary is allowed a chosen-message attack.) The first of these notions is the more natural security requirement while the interest of the second, stronger notion is perhaps more in the implications we discuss below.

These notions of authenticity are by themselves quite disjoint from the notions of privacy; for example, sending the message in the clear with an accompanying (strong) MAC achieves INT-CTXT but no kind of privacy. To make for useful comparisons, we consider each notion of authenticity coupled with IND-CPA, the weakest notion of privacy; namely the notions on which we focus for comparison purposes are $\text{INT-PTXT} \wedge \text{IND-CPA}$ and $\text{INT-CTXT} \wedge \text{IND-CPA}$. (Read “ \wedge ” as “and”.)

Figure 1 shows the graph of relations between these notions and the above-mentioned older ones in the style of [3]. An “implication” $\mathbf{A} \rightarrow \mathbf{B}$ means that every symmetric encryption scheme meeting notion \mathbf{A} also meets notion \mathbf{B} . A “separation” $\mathbf{A} \not\rightarrow \mathbf{B}$ means that there exists a symmetric encryption scheme meeting notion \mathbf{A} but not notion \mathbf{B} . (This under the minimal assumption that some scheme meeting notion \mathbf{A} exists since otherwise the question is moot.) Only a minimal set of relations is explicitly indicated; the relation between any two notions can be derived from the shown ones. (For example, IND-CCA does not imply $\text{INT-CTXT} \wedge \text{IND-CPA}$ because otherwise, by following arrows, we would get $\text{IND-CCA} \rightarrow \text{INT-PTXT} \wedge \text{IND-CPA}$ contradicting a stated separation.) The dotted lines are reminders of existing relations while the numbers annotating the dark lines are pointers to Propositions or Theorems in this paper.

A few points may be worth highlighting. Integrity of ciphertexts—even when coupled only with the weak privacy requirement IND-CPA—emerges as the most powerful notion. Not only does it imply security against chosen-ciphertext attack, but it is strictly stronger than this notion. Non-malleability—whether under chosen-plaintext or

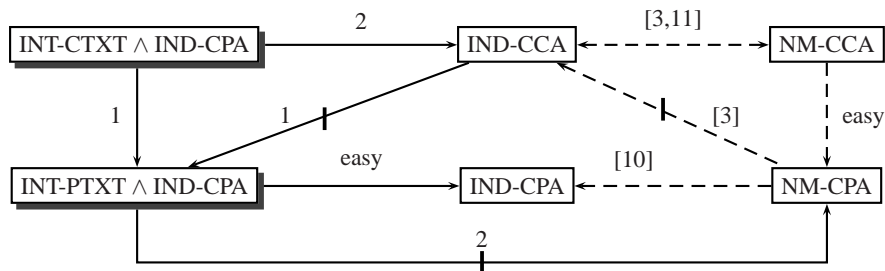


Fig. 1. Relations among notions of symmetric encryption: An arrow denotes an implication while a barred arrow denotes a separation. The full arrows are relations proved in this paper, annotated with the number of the corresponding Proposition or Theorem, while dotted arrows are reminders of existing relations, annotated with citations to the papers establishing them.

chosen-ciphertext attack—does not imply any type of integrity. The intuitive reason is that non-malleability only prevents the generation of ciphertexts whose plaintexts are meaningfully related to those of some challenge ciphertexts, while integrity requires it to be hard to generate ciphertexts of new plaintexts even if these are unrelated to plaintexts underlying any existing ciphertexts. Finally, $\text{INT-PTXT} \wedge \text{IND-CPA}$ does not imply $\text{INT-CTXT} \wedge \text{IND-CPA}$.

1.2 Analysis of Generic Composition

There are many possible ways to design authenticated encryption schemes. We focus in this paper on “generic composition:” simply combine a standard symmetric encryption scheme with a MAC in some way. There are a few possible ways to do it, and our goal is to analyze and compare their security. (The motivation, as we will argue, is that these “obvious” methods, as often the case in practice, remain the most pragmatic from the point of view of performance and security architecture design.)

GENERIC COMPOSITION. Assume we are given a symmetric encryption scheme \mathcal{SE} specified by an encryption algorithm \mathcal{E} and a decryption algorithm \mathcal{D} . (Typically this will be a block cipher mode of operation.) Also assume we are given a message authentication scheme \mathcal{MA} specified by a tagging algorithm \mathcal{T} and a tag verifying algorithm \mathcal{V} and meeting some appropriate notion of unforgeability under chosen-message attack. (Possibilities include the CBC-MAC, HMAC [1], or UMAC [8]). We consider the following methods of “composing” these schemes in order to create an authenticated encryption scheme meeting either $\text{INT-CTXT} \wedge \text{IND-CPA}$ or $\text{INT-PTXT} \wedge \text{IND-CPA}$. We call them “generic” because the algorithms of the authenticated encryption scheme appeal to the given ones as black-boxes only:

Composition Method	Privacy			Integrity	
	IND-CPA	IND-CCA	NM-CPA	INT-PTXT	INT-CTXT
<i>Encrypt-and-MAC plaintext</i>	insecure	insecure	insecure	secure	insecure
<i>MAC-then-encrypt</i>	secure	insecure	insecure	secure	insecure
<i>Encrypt-then-MAC</i>	secure	insecure	insecure	secure	insecure

Fig. 2. Summary of security results for the composed authenticated encryption schemes under the assumption that the given encryption scheme is IND-CPA and the given MAC is weakly unforgeable.

Composition Method	Privacy			Integrity	
	IND-CPA	IND-CCA	NM-CPA	INT-PTXT	INT-CTXT
<i>Encrypt-and-MAC plaintext</i>	insecure	insecure	insecure	secure	insecure
<i>MAC-then-encrypt</i>	secure	insecure	insecure	secure	insecure
<i>Encrypt-then-MAC</i>	secure	secure	secure	secure	secure

Fig. 3. Summary of security results for the composed authenticated encryption schemes under the assumption that the given encryption scheme is IND-CPA and the given MAC is strongly unforgeable.

-
- *Encrypt-and-MAC plaintext*: $\bar{\mathcal{E}}_{K_e, K_m}(M) = \mathcal{E}_{K_e}(M) \| \mathcal{T}_{K_m}(M)$.¹ Namely, encrypt the plaintext and append a MAC of the plaintext. “Decrypt+verify” is performed by first decrypting to get the plaintext and then verifying the tag.
 - *MAC-then-encrypt*: $\bar{\mathcal{E}}_{K_e, K_m}(M) = \mathcal{E}_{K_e}(M \| \mathcal{T}_{K_m}(M))$. Namely, append a MAC to the plaintext and then encrypt them together. “Decrypt+verify” is performed by first decrypting to get the plaintext and candidate tag, and then verifying the tag.
 - *Encrypt-then-MAC*: $\bar{\mathcal{E}}_{K_e, K_m}(M) = C \| \mathcal{T}_{K_m}(C)$ where $C = \mathcal{E}_{K_e}(M)$. Namely, encrypt the plaintext to get a ciphertext C and append a MAC of C . “Decrypt+verify” is performed by first verifying the tag and then decrypting C . This is the method of Internet RFC [16].

Here $\bar{\mathcal{E}}$ is the encryption algorithm of the authenticated encryption scheme while the “decrypt+verify” process specifies a decryption algorithm $\bar{\mathcal{D}}$. The latter will either return a plaintext or a special symbol indicating that it considers the ciphertext unauthentic.

SECURITY RESULTS. Figure 2 and Figure 3 summarize the security results for the three composite authenticated encryption schemes. (We omit NM-CCA since it is equivalent to IND-CCA). Figure 2 shows the results assuming that the base MAC is weakly unforgeable while Figure 3 shows the results assuming that the MAC is strongly unforge-

¹ Here (and everywhere in this paper) “ $\|$ ” denotes an operation that combines several strings into one in such a way that the constituent strings are uniquely recoverable from the final one. (If lengths of all strings are fixed and known, concatenation will serve the purpose.)

able. Weak unforgeability is the standard notion [4]— it should be computationally infeasible for the adversary to find a message-tag pair in which the message is “new,” even after a chosen-message attack. Strong unforgeability requires that it be computationally infeasible for the adversary to find a new message-tag pair even after a chosen-message attack. (The message does not have to be new as long as the output tag was not previously attached to this message by the legitimate parties.) We note that any pseudorandom function is a strongly unforgeable MAC, and most practical MACs seem to be strongly unforgeable. Therefore, analyzing the composition methods under this notion is a realistic and useful approach. Entries in the above tables have the following meaning:

- *Secure*: The composite encryption scheme in question is proven to meet the security requirement in question, assuming only that the component encryption scheme meets IND-CPA and the message authentication scheme is unforgeable under chosen-message attack.
- *Insecure*: There exists *some* IND-CPA secure symmetric encryption and some message authentication scheme unforgeable under chosen-message attack such that the composite scheme based on them does not meet the security requirement in question.

As we can see from Figure 3, the *encrypt-then-MAC* method of [16] is secure from all points of view, making it a good choice for a standard.

The use of a generic composition method secure in the sense above is advantageous from the point of view both of performance and of security architecture. The performance benefit arises from the presence of fast MACs such as HMAC [1] and UMAC [8]. The architectural benefits arise from the stringent notion of security being used. To be secure, the composition must be secure for *all* possible secure instantiations of its constituent primitives. (If it is secure for some instantiations but not others, we declare it insecure.) An application can thus choose a symmetric encryption scheme and a message authentication scheme independently (these are usually already supported by existing security analyses) and then appeal to some fixed and standard composition technique to combine them. No tailored security analysis of the composed scheme is required.

In Section 4 we state formal theorems to support the above claims, providing quantitative bounds for the positive results, and counter-examples with attacks for the negative result. For brevity, we provide theorems and proofs for only the results in Figure 3 (i.e. the strong MAC case).

QUANTITATIVE RESULTS AND COMPARISONS. Above we have discussed our results at a qualitative level. Each result also has a quantitative counterpart; these are what our theorems actually state and prove. These “concrete security” analyses enable a designer to estimate the security of the authenticated encryption scheme in terms of that of its components. All the reductions in this paper are tight, meaning there is little to no loss of security.

1.3 Related Work

The notions IND-CCA, NM-CCA were denoted IND-CCA2 and NM-CCA2, respectively, in [3]. The chosen-ciphertext attacks here are the adaptive kind [18]. Consideration of non-adaptive chosen-ciphertext attacks [17] leads to two more notions, denoted IND-CCA1 and NM-CCA1 by [3], who worked out the relations between six notions of privacy, these two and the four we consider here. (Their results hold for both the asymmetric and the symmetric settings, as mentioned before.) Three additional notions of privacy are considered and related to these six by [14]. In this paper, we have for simplicity avoided consideration of all the possible notions of privacy, focusing instead on what we consider the (four) main ones and their relations to the notions of authenticity. Relations of the remaining notions of privacy to the notions of authenticity considered here can be easily worked out.

Authenticity of an encryption scheme has been understood as a goal by designers for many years. The first formalization of which we are aware is that of [6]. (Early versions of their work date to 1998.) The notion they formalized was INT-CTXT. The formalization of INT-PTXT we use here seems to be new. In independent and concurrent work (both papers were submitted to FSE00) Katz and Yung [15] formalize INT-CTXT plus two other notions of authenticity not considered here. They also observe the implication $\text{INT-CTXT} \wedge \text{IND-CPA} \rightarrow \text{IND-CCA}$.

Generic composition is one of many approaches to the design of authenticated encryption schemes. Two more general approaches are “encryption with redundancy” — append redundancy to the message before encrypting, the latter typically with some block cipher mode of operation— and “encode then encipher” [6] —add randomness and redundancy and then encipher rather than encrypt. As indicated above, attacks have been found on many encrypt with redundancy schemes. Encode then encipher, however, can be proven to work [6] —meaning yields schemes achieving $\text{INT-CTXT} \wedge \text{IND-CPA}$ — but requires a variable-input length pseudorandom permutation, which can be relatively expensive to construct. In addition, there are many specific schemes. One such scheme is the RPC mode of [15] but it is computation and space inefficient compared to the generic composition methods. (Processing an n -block plaintext requires $(1 + c)n$ block cipher computations and results in a ciphertext of this many blocks, where $c \geq 0.3$.) Another scheme is the elegant IACBC mode of Jutla [13] which uses $n + O(\log n)$ block cipher operations to process an n -block plaintext. Implementation and testing would be required to compare its speed with that of generic composition methods that use fast MACs (cf. [1,8]).

Authenticated encryption is not the only approach to achieving security against chosen-ciphertext attacks. Direct approaches yielding more compact schemes have been provided by Desai [9].

2 Definitions

We present definitions for symmetric encryption following [2], first specifying the *syntax* —meaning what kinds of algorithms make up the scheme— and then specifying formal security measures. Associated with each scheme, each notion of security and each adversary is an advantage function that measures the success probability of this

adversary as a function of the security parameter. We define asymptotic notions of security result by asking this function to be negligible for adversaries of time complexity polynomial in the security parameter. Concrete security assessments are made by associating to the scheme another advantage function that for each value of the security parameter and given resources for an adversary returns the maximum, over all adversaries limited to the given resources, of the success probability.

The concrete security assessments are important in practical applications—block cipher based schemes have no associated asymptotics. Hence, we provide concrete security assessments for all positive results (implications or proofs that composition methods meet some notion of security). For simplicity, however, negative results (separations or counter-examples) are phrased in the asymptotic style. (Concrete security statements are, however, easily derived from the proofs.)

SYNTAX OF (SYMMETRIC) ENCRYPTION SCHEMES. A (*symmetric*) *encryption scheme* $\mathcal{SE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ consists of three algorithms. The *randomized key generation* algorithm \mathcal{K} takes input a security parameter $k \in \mathbb{N}$ and returns a key K ; we write $K \xleftarrow{R} \mathcal{K}(k)$. The *encryption* algorithm \mathcal{E} could be randomized or stateful. It takes the key K and a *plaintext* M to return a *ciphertext* C ; we write $C \xleftarrow{R} \mathcal{E}_K(M)$. (If randomized, it flips coins anew on each invocation. If stateful, it uses and then updates a state that is maintained across invocations.) The *decryption* algorithm \mathcal{D} is deterministic and stateless. It takes the key K and a string C to return either the corresponding plaintext M or the symbol \perp ; we write $x \leftarrow \mathcal{D}_K(C)$ where $x \in \{0, 1\}^* \cup \{\perp\}$. We require that $\mathcal{D}_K(\mathcal{E}_K(M)) = M$ for all $M \in \{0, 1\}^*$. An authenticated encryption scheme is syntactically identical to an encryption scheme as defined above; we will use the term only to emphasize cases where we are targeting authenticity goals.

PRIVACY. We measure indistinguishability via the “left-or-right” model of [2]. Define the *left-or-right* oracle $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$, where $b \in \{0, 1\}$, to take input (x_0, x_1) and do the following: if $b = 0$ it computes $C \leftarrow \mathcal{E}_K(x_0)$ and returns C ; else it computes $C \leftarrow \mathcal{E}_K(x_1)$ and returns C . The adversary makes oracle queries of the form (x_0, x_1) consisting of two equal length messages and must guess the bit b . To model chosen-ciphertext attacks we allow the adversary to also have access to a decryption oracle.

Definition 1. (Indistinguishability of a Symmetric Encryption Scheme [2]) Let $\mathcal{SE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme. Let $b \in \{0, 1\}$ and $k \in \mathbb{N}$. Let A_{cpa} be an adversary that has access to the oracle $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ and let A_{cca} be an adversary that has access to the oracles $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ and $\mathcal{D}_K(\cdot)$. Now, we consider the following experiments:

Experiment $\text{Exp}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{ind-cpa-}b}(k)$	Experiment $\text{Exp}_{\mathcal{SE}, A_{\text{cca}}}^{\text{ind-cca-}b}(k)$
$K \xleftarrow{R} \mathcal{K}(k)$	$K \xleftarrow{R} \mathcal{K}(k)$
$x \leftarrow A_{\text{cpa}}^{\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))}(k)$	$x \leftarrow A_{\text{cca}}^{\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b)), \mathcal{D}_K(\cdot)}(k)$
Return x	Return x

Above it is mandated that A_{cca} never queries $\mathcal{D}_K(\cdot)$ on a ciphertext C output by the $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ oracle, and that the two messages queried of $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ always have equal length. We define the *advantages* of the adversaries via

$$\begin{aligned}\mathbf{Adv}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{ind-cpa}}(k) &= \Pr \left[\mathbf{Exp}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{ind-cpa-1}}(k) = 1 \right] - \Pr \left[\mathbf{Exp}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{ind-cpa-0}}(k) = 1 \right] \\ \mathbf{Adv}_{\mathcal{SE}, A_{\text{cca}}}^{\text{ind-cca}}(k) &= \Pr \left[\mathbf{Exp}_{\mathcal{SE}, A_{\text{cca}}}^{\text{ind-cca-1}}(k) = 1 \right] - \Pr \left[\mathbf{Exp}_{\mathcal{SE}, A_{\text{cca}}}^{\text{ind-cca-0}}(k) = 1 \right].\end{aligned}$$

We define the *advantage functions of the scheme* as follows. For any integers t, q_e, q_d, μ ,

$$\begin{aligned}\mathbf{Adv}_{\mathcal{SE}}^{\text{ind-cpa}}(k, t, q_e, \mu) &= \max_{A_{\text{cpa}}} \{ \mathbf{Adv}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{ind-cpa}}(k) \} \\ \mathbf{Adv}_{\mathcal{SE}}^{\text{ind-cca}}(k, t, q_e, q_d, \mu) &= \max_{A_{\text{cca}}} \{ \mathbf{Adv}_{\mathcal{SE}, A_{\text{cca}}}^{\text{ind-cca}}(k) \}\end{aligned}$$

where the maximum is over all $A_{\text{cpa}}, A_{\text{cca}}$ with “time complexity” t , each making at most q_e queries to the $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ oracle, totaling at most μ bits, and, in the case of A_{cca} , also making at most q_d queries to the $\mathcal{D}_K(\cdot)$ oracle. The scheme \mathcal{SE} is said to be *IND-CPA secure* (resp. *IND-CCA secure*) if the function $\mathbf{Adv}_{\mathcal{SE}, A}^{\text{ind-cpa}}(\cdot)$ (resp. $\mathbf{Adv}_{\mathcal{SE}, A}^{\text{ind-cca}}(\cdot)$) is negligible for any adversary A whose time complexity is polynomial in k . ■

The “time complexity” is the worst case total execution time of the experiment, plus the size of the code of the adversary, in some fixed RAM model of computation. We stress that the total execution time of the experiment includes the time of *all* operations in the experiment, including the time for key generation and the computation of answers to oracle queries. Thus, when the time complexity is polynomially bounded, so are all the other parameters. This convention for measuring time complexity and other resources of an adversary is used for all definitions in this paper. The advantage function is the maximum probability that the security of the scheme \mathcal{SE} can be compromised by an adversary using the indicated resources, and is used for concrete security analyses.

We will not use definitions of non-malleability as per [10,3] but instead use the equivalent indistinguishability under parallel chosen-ciphertext attack characterization of [7]. This facilitates our proofs and analyses and also facilitates concrete security measurements. The notation $\mathcal{D}_K(\cdot)$ denotes the algorithm which takes input a vector $\mathbf{c} = (c_1, \dots, c_n)$ of ciphertexts and returns the corresponding vector $\mathbf{p} = (\mathcal{D}_K(c_1), \dots, \mathcal{D}_K(c_n))$ of plaintexts.

Definition 2. (Non-Malleability of a Symmetric Encryption Scheme [7]) Let $\mathcal{SE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme. Let $b \in \{0, 1\}$ and $k \in \mathbb{N}$. Let $A_{\text{cpa}} = (A_{\text{cpa}_1}, A_{\text{cpa}_2})$ be an adversary that has access to the oracle $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ and let $A_{\text{cca}} = (A_{\text{cca}_1}, A_{\text{cca}_2})$ be an adversary that has access to the oracles $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ and $\mathcal{D}_K(\cdot)$. Now, we consider the following experiments:

<p>Experiment $\mathbf{Exp}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{nm-cpa-}b}(k)$</p> <p>$K \xleftarrow{R} \mathcal{K}(k)$</p> <p>$(\mathbf{c}, s) \leftarrow A_{\text{cpa}_1}^{\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))}(k)$</p> <p>$\mathbf{p} \leftarrow \mathcal{D}_K(\mathbf{c})$</p> <p>$x \leftarrow A_{\text{cpa}_2}(\mathbf{p}, \mathbf{c}, s)$</p> <p>Return x</p>	<p>Experiment $\mathbf{Exp}_{\mathcal{SE}, A_{\text{cca}}}^{\text{nm-cca-}b}(k)$</p> <p>$K \xleftarrow{R} \mathcal{K}(k)$</p> <p>$(\mathbf{c}, s) \leftarrow A_{\text{cca}_1}^{\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b)), \mathcal{D}_K(\cdot)}(k)$</p> <p>$\mathbf{p} \leftarrow \mathcal{D}_K(\mathbf{c})$</p> <p>$x \leftarrow A_{\text{cca}_2}(\mathbf{p}, \mathbf{c}, s)$</p> <p>Return x</p>
---	---

Above it is mandated that the vector c output by A_{cpa_1} does not contain any of the ciphertexts output by the $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ oracle, and that the pairs of messages queried of $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ are always of equal length. We define the *advantages* of the adversaries via

$$\begin{aligned}\text{Adv}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{nm-cpa}}(k) &= \Pr \left[\text{Exp}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{nm-cpa-1}}(k) = 1 \right] - \Pr \left[\text{Exp}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{nm-cpa-0}}(k) = 1 \right] \\ \text{Adv}_{\mathcal{SE}, A_{\text{cca}}}^{\text{nm-cca}}(k) &= \Pr \left[\text{Exp}_{\mathcal{SE}, A_{\text{cca}}}^{\text{nm-cca-1}}(k) = 1 \right] - \Pr \left[\text{Exp}_{\mathcal{SE}, A_{\text{cca}}}^{\text{nm-cca-0}}(k) = 1 \right].\end{aligned}$$

We define the *advantage functions of the scheme* as follows. For any integers t, q_e, q_d, μ ,

$$\begin{aligned}\text{Adv}_{\mathcal{SE}}^{\text{nm-cpa}}(k, t, q_e, \mu) &= \max_{A_{\text{cpa}}} \{ \text{Adv}_{\mathcal{SE}, A_{\text{cpa}}}^{\text{nm-cpa}}(k) \} \\ \text{Adv}_{\mathcal{SE}}^{\text{nm-cca}}(k, t, q_e, q_d, \mu) &= \max_{A_{\text{cca}}} \{ \text{Adv}_{\mathcal{SE}, A_{\text{cca}}}^{\text{nm-cca}}(k) \}\end{aligned}$$

where the maximum is over all $A_{\text{cpa}}, A_{\text{cca}}$ with time complexity t , each making at most q_e queries to the $\mathcal{E}_K(\mathcal{LR}(\cdot, \cdot, b))$ oracle, totaling at most μ bits, and, in the case of A_{cca} , also making at most q_d queries to the $\mathcal{D}_K(\cdot)$ oracle. The scheme \mathcal{SE} is said to be *NM-CPA secure* (resp. *NM-CCA secure*) if the function $\text{Adv}_{\mathcal{SE}, A}^{\text{nm-cpa}}(\cdot)$ (resp. $\text{Adv}_{\mathcal{SE}, A}^{\text{nm-cca}}(\cdot)$) is negligible for any adversary A whose time complexity is polynomial in k . ■

INTEGRITY. Now we specify security definitions for integrity (authenticity) of a symmetric encryption scheme $\mathcal{SE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. It is convenient to define an algorithm $\mathcal{D}_K^*(\cdot)$ as follows: If $\mathcal{D}_K(C) \neq \perp$, then return 1 Else return 0. We call this the *verification algorithm* or *verification oracle*. The adversary is allowed a chosen-message attack on the scheme, modeled by giving it access to an encryption oracle $\mathcal{E}_K(\cdot)$. It is successful if it makes the verification oracle accept a ciphertext that was not “legitimately produced.” Different interpretations of the latter give rise to different notions.

Definition 3. (Integrity of an Authenticated Encryption Scheme) Let $\mathcal{SE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme. Let $k \in \mathbb{N}$, and let A_{ptxt} and A_{ctxt} be adversaries each of which has access to two oracles: $\mathcal{E}_K(\cdot)$ and $\mathcal{D}_K^*(\cdot)$. Consider these experiments.

Experiment $\text{Exp}_{\mathcal{SE}, A_{\text{ptxt}}}^{\text{int-ptxt}}(k)$	Experiment $\text{Exp}_{\mathcal{SE}, A_{\text{ctxt}}}^{\text{int-ctxt}}(k)$
$K \xleftarrow{R} \mathcal{K}(k)$	$K \xleftarrow{R} \mathcal{K}(k)$
If $A_{\text{ptxt}}^{\mathcal{E}_K(\cdot), \mathcal{D}_K^*(\cdot)}(k)$ makes a query C to the oracle $\mathcal{D}_K^*(\cdot)$ such that	If $A_{\text{ctxt}}^{\mathcal{E}_K(\cdot), \mathcal{D}_K^*(\cdot)}(k)$ makes a query C to the oracle $\mathcal{D}_K^*(\cdot)$ such that
– $\mathcal{D}_K^*(C)$ returns 1, and	– $\mathcal{D}_K^*(C)$ returns 1, and
– $M \stackrel{\text{def}}{=} \mathcal{D}_K(C)$ was never a query to $\mathcal{E}_K(\cdot)$	– C was never a response of $\mathcal{E}_K(\cdot)$
then return 1 else return 0.	then return 1 else return 0.

We define the *advantages* of the adversaries via

$$\begin{aligned}\text{Adv}_{\mathcal{SE}, A_{\text{ptxt}}}^{\text{int-ptxt}}(k) &= \Pr \left[\text{Exp}_{\mathcal{SE}, A_{\text{ptxt}}}^{\text{int-ptxt}}(k) = 1 \right] \\ \text{Adv}_{\mathcal{SE}, A_{\text{ctxt}}}^{\text{int-ctxt}}(k) &= \Pr \left[\text{Exp}_{\mathcal{SE}, A_{\text{ctxt}}}^{\text{int-ctxt}}(k) = 1 \right]\end{aligned}$$

We define the *advantage functions of the scheme* as follows. For any integers t, q_e, q_d, μ ,

$$\begin{aligned}\mathbf{Adv}_{\mathcal{SE}}^{\text{int-ptxt}}(k, t, q_e, q_d, \mu) &= \max_{A_{\text{ptxt}}} \{ \mathbf{Adv}_{\mathcal{SE}, A_{\text{ptxt}}}^{\text{int-ptxt}}(k) \} \\ \mathbf{Adv}_{\mathcal{SE}}^{\text{int-ctxt}}(k, t, q_e, q_d, \mu) &= \max_{A_{\text{ctxt}}} \{ \mathbf{Adv}_{\mathcal{SE}, A_{\text{ctxt}}}^{\text{int-ctxt}}(k) \}\end{aligned}$$

where the maximum is over all $A_{\text{ptxt}}, A_{\text{ctxt}}$ with time complexity t , each making at most q_e queries to the oracle $\mathcal{E}_K(\cdot)$ and at most q_d queries to $\mathcal{D}_K^*(\cdot)$ such that the sum of the lengths of all oracle queries is at most μ bits. The scheme \mathcal{SE} is said to be *INT-PTXT secure* (resp. *INT-CTXT secure*) if the function $\mathbf{Adv}_{\mathcal{SE}, A}^{\text{int-ptxt}}(\cdot)$ (resp. $\mathbf{Adv}_{\mathcal{SE}, A}^{\text{int-ctxt}}(\cdot)$) is negligible for any adversary A whose time complexity is polynomial in k . ■

MESSAGE AUTHENTICATION SCHEMES. A *message authentication scheme* $\mathcal{MA} = (\mathcal{K}, \mathcal{T}, \mathcal{V})$ consists of three algorithms. The randomized *key generation* algorithm \mathcal{K} takes input a security parameter $k \in \mathbb{N}$ and returns a key K ; we write $K \xleftarrow{R} \mathcal{K}(k)$. The *tagging* algorithm \mathcal{T} could be either randomized or stateful. It takes the key K and a message M to return a tag σ ; we write $\sigma \xleftarrow{R} \mathcal{T}_K(M)$. The *verification* algorithm \mathcal{V} is deterministic. It takes the key K , a message M , and a candidate tag σ for M to return a bit v ; we write $v \leftarrow \mathcal{V}_K(M, \sigma)$. We require that $\mathcal{V}_K(M, \mathcal{T}_K(M)) = 1$ for all $M \in \{0, 1\}^*$. The scheme is said to be deterministic if the tagging algorithm is deterministic and verification is done via tag re-computation. We sometimes call a message authentication scheme a MAC, and also sometimes call the tag σ a MAC.

Security for message authentication considers an adversary F who is allowed a chosen-message attack, modeled by allowing it access to an oracle for $\mathcal{T}_K(\cdot)$. F is “successful” if it can make the verifying oracle $\mathcal{V}_K(\cdot, \cdot)$ accept a pair (M, σ) that was not “legitimately produced.” There are two possible conventions with regard to what “legitimately produced” can mean, leading to two measures of advantage. In the following definition, we use the acronyms WUF-CMA and SUF-CMA respectively for weak and strong unforgeability against chosen-message attacks.

Definition 4. (Message Authentication Scheme Security) Let $\mathcal{MA} = (\mathcal{K}, \mathcal{T}, \mathcal{V})$ be a message authentication scheme. Let $k \in \mathbb{N}$, and let F_w and F_s be adversaries that have access to two oracles: $\mathcal{T}_K(\cdot)$ and $\mathcal{V}_K(\cdot, \cdot)$. Consider the following experiment:

Experiment $\mathbf{Exp}_{\mathcal{MA}, \mathcal{F}_w}^{\text{wuf-cma}}(k)$	Experiment $\mathbf{Exp}_{\mathcal{MA}, \mathcal{F}_s}^{\text{suf-cma}}(k)$
$K \xleftarrow{R} \mathcal{K}(k)$	$K \xleftarrow{R} \mathcal{K}(k)$
If $F_w^{\mathcal{T}_K(\cdot), \mathcal{V}_K(\cdot, \cdot)}(k)$ makes a query (M, σ)	If $F_s^{\mathcal{T}_K(\cdot), \mathcal{V}_K(\cdot, \cdot)}(k)$ makes a query (M, σ)
to the oracle $\mathcal{V}_K(\cdot, \cdot)$ such that	to the oracle $\mathcal{V}_K(\cdot, \cdot)$ such that
– $\mathcal{V}_K(M, \sigma)$ returns 1, and	– $\mathcal{V}_K(M, \sigma)$ returns 1, and
– M was never queried to	– σ was never returned by the
the oracle $\mathcal{T}_K(\cdot)$,	oracle $\mathcal{T}_K(\cdot)$ in response to query M ,
then return 1 else return 0.	then return 1 else return 0.

We define the *advantages* of the forgers via

$$\begin{aligned}\mathbf{Adv}_{\mathcal{MA}, F_w}^{\text{wuf-cma}}(k) &= \Pr \left[\mathbf{Exp}_{\mathcal{MA}, \mathcal{F}_w}^{\text{wuf-cma}}(k) = 1 \right] \\ \mathbf{Adv}_{\mathcal{MA}, F_s}^{\text{suf-cma}}(k) &= \Pr \left[\mathbf{Exp}_{\mathcal{MA}, \mathcal{F}_s}^{\text{suf-cma}}(k) = 1 \right]\end{aligned}$$

We define the *advantage functions of the scheme* as follows. For any integers t, q_t, q_v, μ ,

$$\begin{aligned}\text{Adv}_{\mathcal{MA}}^{\text{wuf-cma}}(k, t, q_t, q_v, \mu) &= \max_{F_w} \{ \text{Adv}_{\mathcal{MA}, F_w}^{\text{wuf-cma}}(k) \} \\ \text{Adv}_{\mathcal{MA}}^{\text{suf-cma}}(k, t, q_t, q_v, \mu) &= \max_{F_s} \{ \text{Adv}_{\mathcal{MA}, F_s}^{\text{suf-cma}}(k) \}\end{aligned}$$

where the maximum is over all F_w, F_s with time complexity t , making at most q_t oracle queries to $\mathcal{T}_K(\cdot)$ and at most q_v oracle queries to $\mathcal{V}_K(\cdot, \cdot)$ such that the sum of the lengths of all oracle queries is at most μ bits. The scheme \mathcal{MA} is said to be *WUF-CMA secure* (resp. *SUF-CMA secure*) if the function $\text{Adv}_{\mathcal{MA}, F}^{\text{wuf-cma}}(\cdot)$ (resp. $\text{Adv}_{\mathcal{MA}, F}^{\text{suf-cma}}(\cdot)$) is negligible for any forger F whose time complexity is polynomial in k . ■

3 Relations among Notions

In this section, we state the formal versions of the results summarized in Figure 1. We begin with the implications and then move to the separations. All proofs are in the full version of this paper [5]. The first implication, below, is a triviality:

Theorem 1. (*INT-CTXT* \rightarrow *INT-PTXT*) Let \mathcal{SE} be an encryption scheme. If \mathcal{SE} is *INT-CTXT secure*, then it is *INT-PTXT secure* as well. Concretely:

$$\text{Adv}_{\mathcal{SE}}^{\text{int-ptxt}}(k, t, q_e, q_d, \mu) \leq \text{Adv}_{\mathcal{SE}}^{\text{int-ctxt}}(k, t, q_e, q_d, \mu) . \blacksquare$$

The next implication is more interesting:

Theorem 2. (*INT-CTXT* \wedge *IND-CPA* \rightarrow *IND-CCA*) Let \mathcal{SE} be an encryption scheme. If \mathcal{SE} is *INT-CTXT secure* and *IND-CPA secure*, then it is *IND-CCA secure*. Concretely:

$$\text{Adv}_{\mathcal{SE}}^{\text{ind-cca}}(k, t, q_e, q_d, \mu) \leq 2 \cdot \text{Adv}_{\mathcal{SE}}^{\text{int-ctxt}}(k, t, q_e, q_d, \mu) + \text{Adv}_{\mathcal{SE}}^{\text{ind-cpa}}(k, t, q_e, \mu) . \blacksquare$$

Next we have the formal statements of the separation results.

Proposition 1. (*IND-CCA* $\not\rightarrow$ *INT-PTXT*) Given a symmetric encryption scheme \mathcal{SE} which is *IND-CCA secure*, we can construct a symmetric encryption scheme \mathcal{SE}' which is also *IND-CCA secure* but is not *INT-PTXT secure*. ■

Proposition 2. (*INT-PTXT* \wedge *IND-CPA* $\not\rightarrow$ *NM-CPA*) Given a symmetric encryption scheme \mathcal{SE} which is both *INT-PTXT secure* and *IND-CPA secure*, we can construct a symmetric encryption scheme \mathcal{SE}' which is also both *INT-PTXT secure* and *IND-CPA secure* but is not *NM-CPA secure*. ■

4 Security of the Composite Schemes

We now present the formal security results for the composite schemes as summarized in Figure 3. The proofs can be found in the full version of this paper [5]. Proofs for the results of Figure 2 are omitted.

Throughout this section, $\mathcal{SE} = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ is a given symmetric encryption scheme which is *IND-CPA secure*, $\mathcal{MA} = (\mathcal{K}_m, \mathcal{T}, \mathcal{V})$ is a given message authentication

scheme which is SUF-CMA secure, and $\overline{SE} = (\overline{K}, \overline{E}, \overline{D})$ is a composite scheme according to one of the three methods we are considering. The presentation below is method by method, and in each case we begin by specifying the method in more detail.

We make the simplifying assumption that \mathcal{D} never returns \perp . It can take any string as input, and the output is always some string. (This is without loss of generality because we can modify \mathcal{D} so that instead of returning \perp it just returns some default message. Security under chosen-plaintext attack is unaffected.) However, \overline{D} can and will return \perp at times, and this is crucial for integrity.

ENCRYPT-AND-MAC PLAINTEXT. The composite scheme is defined as follows:

Algorithm $\overline{K}(k)$	Algorithm $\overline{E}_{\langle K_e, K_m \rangle}(M)$	Algorithm $\overline{D}_{\langle K_e, K_m \rangle}(C)$
$K_e \xleftarrow{R} \mathcal{K}_e(k)$	$C' \leftarrow \mathcal{E}_{K_e}(M)$	Parse C as $C' \parallel \tau$
$K_m \xleftarrow{R} \mathcal{K}_m(k)$	$\tau \leftarrow \mathcal{T}_{K_m}(M)$	$M \leftarrow \mathcal{D}_{K_e}(C')$
Return $\langle K_e, K_m \rangle$	$C \leftarrow C' \parallel \tau$	$v \leftarrow \mathcal{V}_{K_m}(M, \tau)$
	Return C	If $v = 1$, return M
		else return \perp .

This composition method does not preserve privacy because the MAC could reveal information about the plaintext.

Proposition 3. (Encrypt-and-MAC plaintext method is not IND-CPA secure) *Given a IND-CPA secure symmetric encryption scheme SE and a SUF-CMA secure message authentication scheme MA , we can construct a message authentication scheme MA' such that MA' is SUF-CMA secure, but the composite scheme \overline{SE} formed by the encrypt-and-MAC plaintext composition method based on SE and MA' is not IND-CPA secure. ■*

Since both IND-CCA and NM-CPA imply IND-CPA, this means that this composition method is also *neither* IND-CCA *nor* NM-CPA secure.

The *encrypt-and-MAC plaintext* composition method, however, inherits the integrity of the MAC in a direct way:

Theorem 3. (Encrypt-and-MAC plaintext method is INT-PTXT secure) *Let SE be a symmetric encryption scheme, let MA be a message authentication scheme, and let \overline{SE} be the encryption scheme obtained from SE and MA via the encrypt-and-MAC plaintext composition method. Then, if MA is SUF-CMA secure, then \overline{SE} is INT-PTXT secure. Concretely:*

$$\text{Adv}_{\overline{SE}}^{\text{int-ptxt}}(k, t, q_e, q_d, \mu) \leq \text{Adv}_{MA}^{\text{suf-cma}}(k, t, q_e, q_d, \mu) . \blacksquare$$

However, this composition method fails in general to provide integrity of ciphertexts. This is because there are secure encryption schemes with the property that a ciphertext can be modified without changing its decryption. When such an encryption scheme is used as the base symmetric encryption scheme, an adversary can query the encryption oracle, modify part of the response, and still submit the result to the verification oracle as a valid ciphertext. The following proposition states this result.

Proposition 4. (Encrypt-and-MAC plaintext method is not INT-CTXT secure) *Given a IND-CPA secure symmetric encryption scheme SE and a SUF-CMA secure message*

authentication scheme \mathcal{MA} , we can construct a symmetric encryption scheme \mathcal{SE}' such that \mathcal{SE}' is IND-CPA secure, but the composite scheme $\overline{\mathcal{SE}}$ formed by the encrypt-and-MAC plaintext composition method based on \mathcal{SE}' and \mathcal{MA} is not INT-CTXT secure. ■

MAC-THEN-ENCRYPT. The composite scheme is defined as follows:

Algorithm $\overline{\mathcal{K}}(k)$ $K_e \xleftarrow{R} \mathcal{K}_e(k)$ $K_m \xleftarrow{R} \mathcal{K}_m(k)$ Return $\langle K_e, K_m \rangle$	Algorithm $\overline{\mathcal{E}}_{\langle K_e, K_m \rangle}(M)$ $\tau \leftarrow \mathcal{T}_{K_m}(M)$ $C \leftarrow \mathcal{E}_{K_e}(M \parallel \tau)$ Return C	Algorithm $\overline{\mathcal{D}}_{\langle K_e, K_m \rangle}(C)$ $M' \leftarrow \mathcal{D}_{K_e}(C)$ Parse M' as $M \parallel \tau$ $v \leftarrow \mathcal{V}_{K_m}(M, \tau)$ If $v = 1$, return M else return \perp .
--	--	---

The MAC-then-encrypt composition method preserves both privacy against chosen-plaintext attack and integrity of plaintexts, as stated in the following theorem.

Theorem 4. (MAC-then-encrypt method is both INT-PTXT and IND-CPA secure) Let \mathcal{MA} be a message authentication scheme, and let \mathcal{SE} be a symmetric encryption scheme secure against chosen-plaintext attacks. Let $\overline{\mathcal{SE}}$ be the encryption scheme obtained from \mathcal{SE} and \mathcal{MA} via the MAC-then-encrypt composition method. Then, if \mathcal{MA} is SUF-CMA secure, then $\overline{\mathcal{SE}}$ is INT-PTXT secure. Furthermore, if \mathcal{SE} is IND-CPA secure, then so is $\overline{\mathcal{SE}}$. Concretely:

$$\begin{aligned} \text{Adv}_{\overline{\mathcal{SE}}}^{\text{int-ptxt}}(k, t_i, q_e, q_d, \mu_i) &\leq \text{Adv}_{\mathcal{MA}}^{\text{suf-cma}}(k, t_i, q_e, q_d, \mu_i) \\ \text{Adv}_{\overline{\mathcal{SE}}}^{\text{ind-cpa}}(k, t_p, q, \mu_p) &\leq \text{Adv}_{\mathcal{SE}}^{\text{ind-cpa}}(k, t_p, q, \mu_p) . \blacksquare \end{aligned}$$

However, the base encryption scheme might be malleable, and this will be inherited by the composite scheme.

Proposition 5. (MAC-then-encrypt method is not NM-CPA secure) Given a IND-CPA secure symmetric encryption scheme \mathcal{SE} and a SUF-CMA secure message authentication scheme \mathcal{MA} , we can construct a symmetric encryption scheme \mathcal{SE}' such that \mathcal{SE}' is IND-CPA secure, but the composite scheme $\overline{\mathcal{SE}}$ formed by the MAC-then-encrypt composition method based on \mathcal{SE}' and \mathcal{MA} is not NM-CPA secure. ■

Since IND-CCA implies NM-CPA, this composition method is also *not* IND-CCA secure. Furthermore, the fact that it is IND-CPA secure but not NM-CPA secure implies that it is not INT-CTXT secure.

ENCRYPT-THEN-MAC. The composite scheme is defined as follows:

Algorithm $\overline{\mathcal{K}}(k)$ $K_e \xleftarrow{R} \mathcal{K}_e(k)$ $K_m \xleftarrow{R} \mathcal{K}_m(k)$ Return $\langle K_e, K_m \rangle$	Algorithm $\overline{\mathcal{E}}_{\langle K_e, K_m \rangle}(M)$ $C' \leftarrow \mathcal{E}_{K_e}(M)$ $\tau' \leftarrow \mathcal{T}_{K_m}(C')$ $C \leftarrow C' \parallel \tau'$ Return C	Algorithm $\overline{\mathcal{D}}_{\langle K_e, K_m \rangle}(C)$ Parse C as $C' \parallel \tau'$ $M \leftarrow \mathcal{D}_{K_e}(C')$ $v \leftarrow \mathcal{V}_{K_m}(C', \tau')$ If $v = 1$, return M else return \perp .
--	---	--

The following theorem implies that the encrypt-then-MAC composition method is IND-CPA, IND-CCA, NM-CPA, INT-PTXT and INT-CTXT secure.

Theorem 5. (*Encrypt-then-MAC method is INT-CTXT, IND-CPA, and IND-CCA secure*) Let \mathcal{SE} be a symmetric encryption scheme, and let \mathcal{MA} be a message authentication scheme. Let $\overline{\mathcal{SE}}$ be the authenticated encryption scheme obtained from \mathcal{SE} and \mathcal{MA} via the encrypt-then-MAC composition method. Then, if \mathcal{MA} is *SUF-CMA* secure, then $\overline{\mathcal{SE}}$ is *INT-CTXT* secure. If \mathcal{SE} is *IND-CPA* secure, then so is $\overline{\mathcal{SE}}$. And if we have both of the previous conditions, then $\overline{\mathcal{SE}}$ is *IND-CCA* secure. Concretely:

$$\begin{aligned}\text{Adv}_{\overline{\mathcal{SE}}}^{\text{int-ctxt}}(k, t_2, q_2, q'_2, \mu_2) &\leq \text{Adv}_{\mathcal{MA}}^{\text{suf-cma}}(k, t_2, q_2, q'_2, \mu_2) \\ \text{Adv}_{\overline{\mathcal{SE}}}^{\text{ind-cpa}}(k, t_3, q_3, \mu_3) &\leq \text{Adv}_{\mathcal{SE}}^{\text{ind-cpa}}(k, t_3, q_3, \mu_3)\end{aligned}$$

and

$$\begin{aligned}\text{Adv}_{\overline{\mathcal{SE}}}^{\text{ind-cca}}(k, t_4, q_4, q'_4, \mu_4) &\leq \\ 2 \cdot \text{Adv}_{\mathcal{MA}}^{\text{suf-cma}}(k, t_4, q_4, q'_4, \mu_4) &+ \text{Adv}_{\mathcal{SE}}^{\text{ind-cpa}}(k, t_4, q_4, \mu_4) \cdot \blacksquare\end{aligned}$$

Acknowledgments

The authors are supported in part by a 1996 Packard Foundation Fellowship in Science and Engineering and NSF CAREER Award CCR-9624439.

References

1. M. BELLARE, R. CANETTI AND H. KRAWCZYK, “Keying hash functions for message authentication,” *Advances in Cryptology – Crypto ’96*, LNCS Vol. 1109, N. Kobitz ed., Springer-Verlag, 1996.
2. M. BELLARE, A. DESAI, E. JOKIPII AND P. ROGAWAY, “A concrete security treatment of symmetric encryption: Analysis of the DES modes of operation,” *Proc. of the 38th IEEE FOCS*, IEEE, 1997.
3. M. BELLARE, A. DESAI, D. POINTCHEVAL AND P. ROGAWAY, “Relations among notions of security for public-key encryption schemes,” *Advances in Cryptology – Crypto ’98*, LNCS Vol. 1462, H. Krawczyk ed., Springer-Verlag, 1998.
4. M. BELLARE, J. KILIAN, P. ROGAWAY, “The security of the cipher block chaining message authentication code,” *Advances in Cryptology – Crypto ’94*, LNCS Vol. 839, Y. Desmedt ed., Springer-Verlag, 1994.
5. M. BELLARE, C. NAMPREMPRE, “Authenticated Encryption: Relations among notions and analysis of the generic composition paradigm,” Full version of this paper, available via <http://www-cse.ucsd.edu/users/mihir>.
6. M. BELLARE AND P. ROGAWAY, “Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient cryptography,” *Advances in Cryptology – ASIACRYPT ’00*, LNCS Vol. ??, T. Okamoto ed., Springer-Verlag, 2000.
7. M. BELLARE AND A. SAHAI, “Non-Malleable Encryption: Equivalence between Two Notions, and an Indistinguishability-Based Characterization,” *Advances in Cryptology – Crypto ’99*, LNCS Vol. 1666, M. Wiener ed., Springer-Verlag, 1999.
8. J. BLACK, S. HALEVI, H. KRAWCZYK, T. KROVETZ AND P. ROGAWAY, “UMAC: Fast and secure message authentication,” *Advances in Cryptology – Crypto ’99*, LNCS Vol. 1666, M. Wiener ed., Springer-Verlag, 1999.

9. A. DESAI, "New paradigms for constructing symmetric encryption schemes secure against chosen ciphertext attack," *Advances in Cryptology – Crypto '00*, LNCS Vol. 1880, M. Bellare ed., Springer-Verlag, 2000.
10. D. DOLEV, C. DWORK, AND M. NAOR, "Non-malleable cryptography," *Proc. of the 23rd ACM STOC*, ACM, 1991.
11. D. DOLEV, C. DWORK, AND M. NAOR, "Non-malleable cryptography," to appear in *SIAM J. Comput.*
12. S. GOLDWASSER AND S. MICALI, "Probabilistic encryption," *Journal of Computer and System Science*, Vol. 28, 1984, pp. 270-299.
13. C. JUTLA, "Encryption modes with almost free message integrity," Report 2000/039, *Cryptology ePrint Archive*, <http://eprint.iacr.org/>, August 2000.
14. J. KATZ AND M. YUNG, "Complete characterization of security notions for probabilistic private-key encryption," *Proc. of the 32nd ACM STOC*, ACM, 2000.
15. J. KATZ AND M. YUNG, "Unforgeable Encryption and Adaptively Secure Modes of Operation," *Fast Software Encryption '00*, LNCS Vol. ??, B. Schneier ed., Springer-Verlag, 2000.
16. S. KENT AND R. ATKINSON, "IP Encapsulating Security Payload (ESP)," Request for Comments 2406, November 1998.
17. M. NAOR AND M. YUNG, "Public-key cryptosystems provably secure against chosen ciphertext attacks," *Proc. of the 22nd ACM STOC*, ACM, 1990.
18. C. RACKOFF AND D. SIMON, "Non-Interactive zero-knowledge proof of knowledge and chosen ciphertext attack," *Advances in Cryptology – Crypto '91*, LNCS Vol. 576, J. Feigenbaum ed., Springer-Verlag, 1991.