

Exercise (Authenticate+Encrypt=NM-CPA). Let $\mathfrak{C} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a IND-CPA secure symmetric encryption scheme and let h be a secure message authentication code with the appropriate message and key domains. Show that the paradigm first authenticate and then encrypt


<p>Key-Gen</p> $\left[\begin{array}{l} \text{sk} \leftarrow \text{Gen} \\ k \leftarrow \mathcal{K} \\ \text{return} (\text{sk}, k) \end{array} \right.$	<p>Auth-Enc_{sk,k}(m)</p> $\left[\begin{array}{l} t \leftarrow h(m, k) \\ \text{return Enc}_{\text{sk}}(m, t) \end{array} \right.$	<p>Auth-Dec_{sk,k}(c)</p> $\left[\begin{array}{l} (m, t) \leftarrow \text{Dec}_{\text{sk}}(c) \\ \text{if } t \neq h(m, k) \text{ then return } \perp \\ \text{else return } m \end{array} \right.$
---	--	---

assures non-malleability under chosen message attacks (NM-CPA security).

Solution. Let the encryption scheme \mathfrak{C} be (t_1, ε_1) -IND-CPA secure and let h be (t_2, ε_2) -secure message authentication code. Let us now consider the advantage of an adversary \mathcal{A} against NM-CPA games:

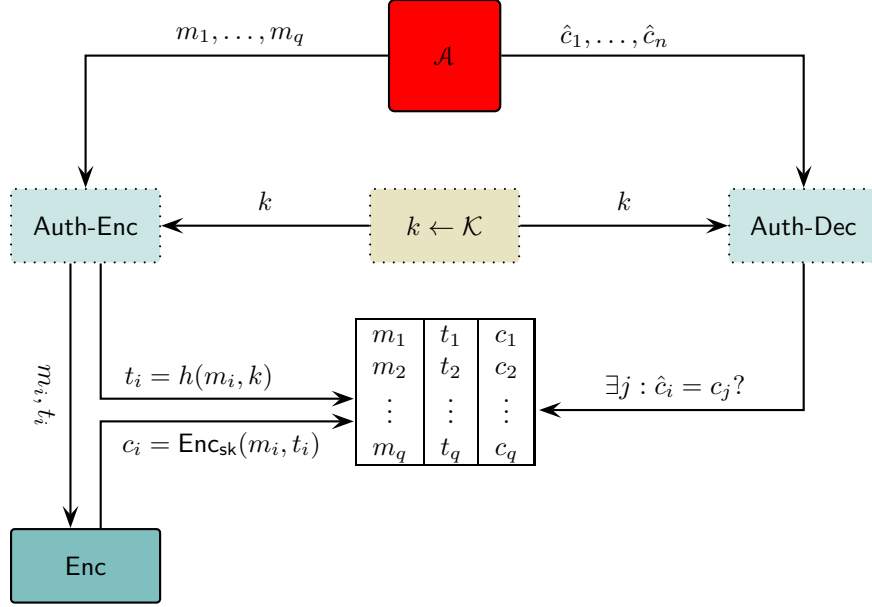
<p>$\mathcal{G}_0^{\mathcal{A}}$</p> $\left[\begin{array}{l} k \leftarrow \mathcal{K} \\ \text{sk} \leftarrow \text{Gen} \\ \mathcal{M}_0 \leftarrow \mathcal{A}^{\text{Auth-Enc}_{\text{sk},k}(\cdot)}(\text{pk}) \\ m \leftarrow \mathcal{M}_0, \bar{m} \leftarrow \mathcal{M}_0 \\ c \leftarrow \text{Enc}_{\text{sk},k}(m) \\ (\pi, \hat{c}_1, \dots, \hat{c}_n) \leftarrow \mathcal{A}^{\text{Auth-Enc}_{\text{sk},k}(\cdot)}(c) \\ \text{if } c \in \{\hat{c}_1, \dots, \hat{c}_n\} \text{ then return } 0 \\ \text{return } \pi(m, \text{Auth-Dec}_{\text{sk},k}(\hat{c}_1), \dots) \end{array} \right.$	<p>$\mathcal{G}_1^{\mathcal{A}}$</p> $\left[\begin{array}{l} k \leftarrow \mathcal{K} \\ \text{sk} \leftarrow \text{Gen} \\ \mathcal{M}_0 \leftarrow \mathcal{A}^{\text{Auth-Enc}_{\text{sk},k}(\cdot)}(\text{pk}) \\ m \leftarrow \mathcal{M}_0, \bar{m} \leftarrow \mathcal{M}_0 \\ \pi, \bar{c} \leftarrow \text{Auth-Enc}_{\text{sk},k}(\bar{m}) \\ (\pi, \hat{c}_1, \dots, \hat{c}_n) \leftarrow \mathcal{A}^{\text{Auth-Enc}_{\text{sk},k}(\cdot)}(\bar{c}) \\ \text{if } c \in \{\hat{c}_1, \dots, \hat{c}_n\} \text{ then return } 0 \\ \text{return } \pi(m, \text{Auth-Dec}_{\text{sk},k}(\hat{c}_1), \dots) \end{array} \right.$
---	---

Although the adversary \mathcal{A} does not make the explicit call to the decryption oracle, the form of the security game implies that \mathcal{A} is entitled to a single decryption call, after which the predicate π determines the output of the game. As a result, the following interaction diagram captures interactions in the game.



As our aim is to concert the NM-CPA adversary \mathcal{A} to an IND-CPA adversary \mathcal{B} , we must be able to perform the decryption operation inside \mathcal{B} without an access to the decryption oracle. At first glance, it seems an impossible task. Fortunately, the task is easy when the ciphertext c_i is legitimate ciphertext that was obtained by \mathcal{A} as a reply to an encryption query m_j . Then we know that $\text{Dec}(c_i) = m_j$.

We will exploit this fact by constructing an interface that stores all encryption queries and corresponding replies as a database, which is later used for mimicking replies for the decryption calls. The schematics of this construction is illustrated below. First, the key of the message authentication code is set up by picking $k \leftarrow \mathcal{K}$. After that \mathcal{B} can simulate $\text{Auth-Enc}_{\text{sk},k}(m_i)$ oracle calls by computing the authentication tag $t_i \leftarrow h(m_i, k)$ and submitting the pair (m_i, t_i) to the encryption oracle $\text{Enc}_{\text{sk}}(\cdot)$. Let c_i be the corresponding reply. Then the triple (m_i, t_i, c_i) is stored in the encryption table. To decrypt $\hat{c}_1, \dots, \hat{c}_n$, the adversary \mathcal{B} uses the encryption table. For each \hat{c}_i , it tries to find matching c_j in the table and then use m_j as the decryption \hat{m}_i . If some \hat{c}_i is missing from the table, the \mathcal{B} halts with zero. Otherwise \mathcal{B} returns the value of $\pi(m, \hat{m}_1, \dots, \hat{m}_n)$



Let LegitEnc denote the event that $\hat{c}_1, \dots, \hat{c}_n$ are all generated by querying the $\text{Auth-Enc}_{\text{sk},k}(\cdot)$ oracle. Then the event $\neg \text{LegitEnc}$ does not automatically mean that the abrupt halt by \mathcal{B} is different from the actual evaluation of the predicate π in the non-malleability games. In those games, the game is also halted if the decryption of $\hat{c}_1, \dots, \hat{c}_n$ fails. More precisely, let GoodDec denote the event that \mathcal{A} produces a decipherable vector of ciphertexts $\hat{c}_1, \dots, \hat{c}_n$, i.e., $\text{Auth-Dec}_{\text{sk},k}(\hat{c}_i) \neq \perp$ for all $i \in \{1, \dots, n\}$. Then we can decompose the success probabilities as follows:

$$\begin{aligned} \Pr [\mathcal{G}_i^{\mathcal{A}} = 1] &= \Pr [\mathcal{G}_i^{\mathcal{A}} = 1 \wedge \text{LegitEnc}] \\ &\quad + \Pr [\mathcal{G}_i^{\mathcal{A}} = 1 \wedge \neg \text{LegitEnc} \wedge \neg \text{GoodDec}] + \Pr [\mathcal{G}_i^{\mathcal{A}} = 1 \wedge \neg \text{LegitEnc} \wedge \text{GoodDec}] . \end{aligned}$$

We can similarly decompose the success probabilities of \mathcal{B} who plays against IND-CPA games \mathcal{Q}_0 and \mathcal{Q}_1 :

$$\begin{aligned} \Pr [\mathcal{Q}_i^{\mathcal{B}} = 1] &= \Pr [\mathcal{Q}_i^{\mathcal{B}} = 1 \wedge \text{LegitEnc}] \\ &\quad + \Pr [\mathcal{Q}_i^{\mathcal{B}} = 1 \wedge \neg \text{LegitEnc} \wedge \neg \text{GoodDec}] + \Pr [\mathcal{Q}_i^{\mathcal{B}} = 1 \wedge \neg \text{LegitEnc} \wedge \text{GoodDec}] , \end{aligned}$$

as the substitution of \mathcal{B} leads to the game that is identical except for the last line. That is, events LegitEnc and GoodDec are well defined in games $\mathcal{Q}_0^{\mathcal{B}}$ and $\mathcal{Q}_1^{\mathcal{B}}$. By the construction

$$\Pr [\mathcal{G}_i^{\mathcal{A}} = 1 \wedge \text{LegitEnc}] = \Pr [\mathcal{Q}_i^{\mathcal{B}} = 1 \wedge \text{LegitEnc}]$$

as \mathcal{B} obtains correct decryption values and thus the execution of the last line is identical. Analogously,

$$\Pr [\mathcal{G}_i^{\mathcal{A}} = 1 \wedge \neg \text{LegitEnc} \wedge \neg \text{GoodDec}] = 0 = \Pr [\mathcal{Q}_i^{\mathcal{B}} = 1 \wedge \neg \text{LegitEnc} \wedge \neg \text{GoodDec}] ,$$

since in both games decryption failure causes the abrupt halting. Finally, we also know that

$$\Pr [Q_0^{\mathcal{B}} = 1 \wedge \neg \text{LegitEnc} \wedge \text{GoodDec}] = \Pr [Q_1^{\mathcal{B}} = 1 \wedge \neg \text{LegitEnc} \wedge \text{GoodDec}] ,$$

since \mathcal{B} halts abruptly for non-legitimately generated ciphertexts. Consequently,

$$\begin{aligned} \text{Adv}_{\mathcal{E}}^{\text{ind-cpa}}(\mathcal{B}) &= |\Pr [Q_0^{\mathcal{B}} = 1] - \Pr [Q_1^{\mathcal{B}} = 1]| \\ &= |\Pr [Q_0^{\mathcal{B}} = 1 \wedge \text{LegitEnc}] - \Pr [Q_1^{\mathcal{B}} = 1 \wedge \text{LegitEnc}]| . \end{aligned}$$

As for any event E that is defined and has identical probability in both games

$$|\Pr [\mathcal{G}_0^{\mathcal{A}} = 1 \wedge E] - \Pr [\mathcal{G}_1^{\mathcal{A}} = 1 \wedge E]| \leq \Pr [E]$$

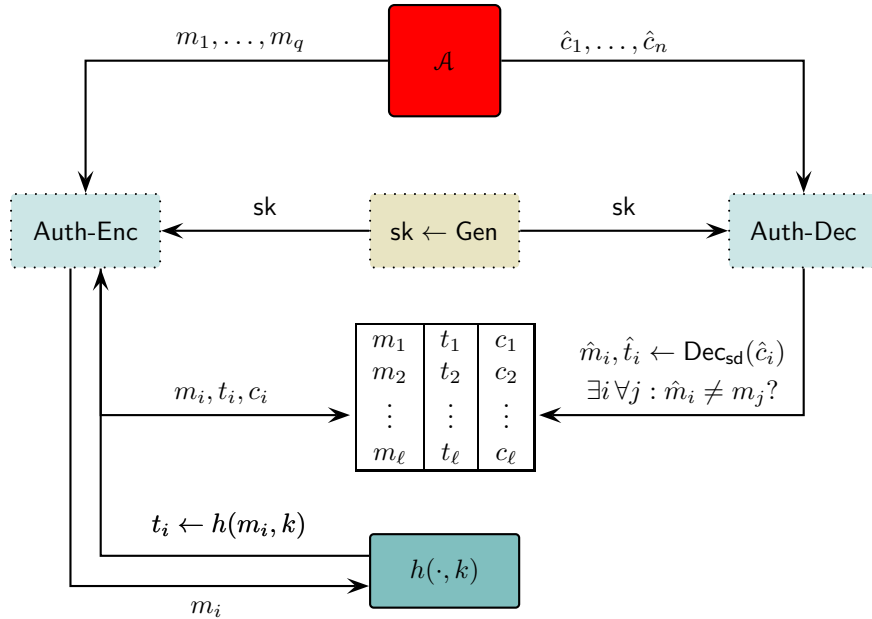
we get by similar reasoning that

$$\begin{aligned} \text{Adv}_{\mathcal{E}}^{\text{nm-cpa}}(\mathcal{A}) &= |\Pr [\mathcal{G}_0^{\mathcal{A}} = 1] - \Pr [\mathcal{G}_1^{\mathcal{A}} = 1]| \\ &\leq |\Pr [\mathcal{G}_0^{\mathcal{A}} = 1 \wedge \text{LegitEnc}] - \Pr [\mathcal{G}_1^{\mathcal{A}} = 1 \wedge \text{LegitEnc}]| + \Pr [\neg \text{LegitEnc} \wedge \text{GoodDec}] . \end{aligned}$$

The latter is equivalent to the inequality

$$\text{Adv}_{\mathcal{E}}^{\text{nm-cpa}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{E}}^{\text{ind-cpa}}(\mathcal{B}) + \Pr [\neg \text{LegitEnc} \wedge \text{GoodDec}] .$$

To complete the proof, we need to assess the second term. Note that if the adversary \mathcal{A} can create a decipherable vector $\hat{c}_1, \dots, \hat{c}_n$ without computing all ciphertexts legitimately, then there must exist \hat{c}_i for which the authentication tag t_i was not computed by the encryption oracle. This leads to new adversary construction \mathcal{C} that simulates \mathcal{G}_0 to the \mathcal{A} in order to get valid MAC forgery.



In this construction, the secret key sk is generated by \mathcal{C} who simulates encryption and decryption queries for the \mathcal{A} . For that, \mathcal{C} must occasionally query hash values from the hash oracle $h(\cdot, k)$ present in the security game defining security of message authentication codes. The adversary \mathcal{C} runs the execution until \mathcal{A} finishes. After that \mathcal{C} decrypts all ciphertexts $\hat{c}_1, \dots, \hat{c}_n$ and halts if the event LegitEnc occurs. Otherwise, there must

exist i such that \hat{m}_i is not queried from the hashing oracle $h(\cdot, k)$. The pair \hat{m}_i, \hat{t}_i will be the output of \mathcal{C} . It is easy to see that if the event **GoodDec** occurs then $\hat{t}_i = h(\hat{m}_i, k)$ and thus \mathcal{C} wins the MAC game. Consequently, we have proven that

$$\Pr[\neg\text{LegitEnc} \wedge \text{GoodDec}] \leq \text{Adv}_h^{\text{mac}}(\mathcal{C}) .$$

Careful analysis shows that the overhead in the constructions is $\Theta(q \log q + n \log q)$. Thus for sufficiently small running time of \mathcal{A} both constructions run in t -time and thus (t, ε_1) -IND-CPA security and (t, ε_2) -security of the message authentication code h assures that

$$\text{Adv}_{\mathcal{C}}^{\text{nm-cpa}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{C}}^{\text{ind-cpa}}(\mathcal{B}) + \text{Adv}_h^{\text{mac}}(\mathcal{C}) \leq \varepsilon_1 + \varepsilon_2 .$$