

# Lecture 8: Authenticated Encryption and CCA Security

Notes by Yael Kalai

MIT - 6.5620

Lecture 8 (September 29, 2025)

**Warning:** This document is a rough draft, so it may contain bugs. Please feel free to email me with corrections.

## Recap

- Last class we proved the security of the GGM construction, which constructs a PRF from any PRG.
- We defined the notion of a message authentication code (MAC), and constructed a MAC from any PRF.

The construction is extremely simple: Given any  $\text{PRF} : \mathcal{K}_\lambda \times \mathcal{M}_\lambda \rightarrow \{0,1\}^\lambda$ , let

$$\text{MAC}(k, m) := \text{PRF}(k, m).$$

## Today

- Prove the security of the MAC construction.
- Construct an Authenticated Encryption scheme.
- Define CCA-secure encryption.
- Prove that an authenticated encryption is CCA-secure.

## Security of the MAC Construction

**Theorem 1.** *The above MAC is existentially unforgeable against adaptive chosen message attacks.*

*Proof.* Fix any poly-size  $\mathcal{A}$ . By the security of the PRF it holds that there exists a negligible  $\mu$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\begin{aligned} \Pr[\mathcal{A}^{\text{MAC}(k, \cdot)}(1^\lambda) = (m^*, \tau^*) : \tau^* = \text{MAC}(k, m^*) \wedge m^* \notin Q] \leq \\ \Pr[\mathcal{A}^{R_\lambda(\cdot)}(1^\lambda) = (m^*, \tau^*) : \tau^* = R_\lambda(m^*) \wedge m^* \notin Q] + \mu(\lambda), \end{aligned}$$

where  $R_\lambda : \mathcal{M}_\lambda \rightarrow \{0,1\}^\lambda$  is a truly random function. It is easy to see that

$$\Pr[\mathcal{A}^{R_\lambda(\cdot)}(1^\lambda) = (m^*, \tau^*) : \tau^* = R_\lambda(m^*) \wedge m^* \notin Q] = 2^{-\lambda}.$$

Thus, we conclude that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr[\mathcal{A}^{\text{MAC}(k,\cdot)}(1^\lambda) = (m^*, \tau^*) : \tau^* = \text{MAC}(k, m^*) \wedge m^* \notin Q] \leq \mu(\lambda) + 2^{-\lambda}$$

which is negligible. □

### MACing Long Messages

Suppose we are given a PRF that takes as input messages in  $\{0,1\}^n$  and we want to MAC messages in  $\{0,1\}^{2n}$ ? In particular, suppose we want to MAC the message  $m_1 || m_2$  where  $m_1, m_2 \in \{0,1\}^n$ . It is tempting to let

$$\text{MAC}(k, m_1 || m_2) = \text{MAC}(k, m_1) || \text{MAC}(k, m_2).$$

However, this will not be secure, since the adversary can “mix-and-match”. Specifically, given

$$\text{MAC}(k, m_1 || m_2) = \text{MAC}(k, m_1) || \text{MAC}(k, m_2)$$

the adversary can compute

$$\text{MAC}(k, m_2 || m_1) = \text{MAC}(k, m_2) || \text{MAC}(k, m_1).$$

Moreover, given

$$\text{MAC}(k, m_1 || m_2) = \text{MAC}(k, m_1) || \text{MAC}(k, m_2)$$

and

$$\text{MAC}(k, m'_1 || m'_2) = \text{MAC}(k, m'_1) || \text{MAC}(k, m'_2),$$

the adversary can generate

$$\text{MAC}(k, m_1 || m'_2) = \text{MAC}(k, m_1) || \text{MAC}(k, m'_2).$$

*Remark.* Note that this is in contrast to CPA secure encryption, where there concatenation preserves security!

So what if we want to MAC a long message? Well, there are several ways to solve this problem. One way is to construct a PRF for messages of arbitrary length by making the GGM tree deeper. Recall that the depth of the tree determines the length of the message. We will talk about other ways in the next lecture.

### Authenticated Encryption

Often in application we want both secrecy and authenticity! So, what should we do? Thankfully, we already constructed a CPA-secure encryption scheme  $(\text{Enc}, \text{Dec})$ , and we also constructed a MAC. But how do we put them together?

*Take 1:*

1. Alice and Bob first share a random secret key  $k \leftarrow \{0, 1\}^\lambda$ .
2. Every time Alice sends a message  $m \in \mathcal{M}_\lambda$  to Bob, she should encrypt her message
 
$$\text{ct} \leftarrow \text{Enc}(k, m).$$
3. Then, since she wishes to send this ciphertext in an authenticated way, she computes  $\tau = \text{MAC}(k, \text{ct})$ .
4. She sends  $(\text{ct}, \tau)$ .

Upon receiving  $(\text{ct}, \tau)$ , Bob first checks that indeed  $\tau = \text{MAC}(k, \text{ct})$ , and if this is not the case it discards this message! If the check passes, then he knows that this ciphertext was indeed sent by Alice, and then he decrypts it to retrieve the message  $m = \text{Dec}(k, \text{ct})$ . Is this secure? Not necessarily!

*Lesson:* Never use the same secret key for different applications! This can leak the secret key entirely! The issue is that both a MAC and a ciphertext may leak information about the secret key, and these leakages may not play nicely with each other!

The correct way to communicate securely, is to share *two* secret keys, one for the encryption and one for the MAC.

*Our authenticated encryption scheme.* Given a CPA-secure encryption  $(\text{Enc}, \text{Dec})$  and given a MAC scheme  $\text{MAC}$ , we define an authenticated encryption scheme, denoted by  $(\text{Enc}', \text{Dec}')$ , as follows:

- First, Alice and Bob share *two* random secret key  $k_{\text{Enc}}, k_{\text{MAC}} \leftarrow \{0, 1\}^\lambda$ .
- $\text{Enc}'((k_{\text{MAC}}, k_{\text{Enc}}), m)$ :
  1. Compute  $c \leftarrow \text{Enc}(k_{\text{Enc}}, m)$ .
  2. Compute  $\tau = \text{MAC}(k_{\text{MAC}}, c)$ .
  3. Output  $\text{ct} = (c, \tau)$ .
- $\text{Dec}'((k_{\text{MAC}}, k_{\text{Enc}}), \text{ct})$ :
  1. Parse  $\text{ct} = (c, \tau)$ .

2. Check if  $\tau = \text{MAC}(k_{\text{MAC}}, c)$ . If this is not the case then output  $\perp$ .
3. Output  $m = \text{Dec}(k_{\text{Enc}}, c)$

*Question:* Why do we first encrypt and then MAC? Why don't we first MAC and then encrypt – namely, send  $\text{Enc}(k_{\text{Enc}}, (m, \text{MAC}(K_{\text{MAC}}, m)))$ ?

The reason is that if we Encrypt-then-Mac (as is done above) then our resulting scheme is secure against *adaptive chosen ciphertext attacks* (CCA), defined below. CCA-security is the golden standard security notion.

**Definition 2.** An encryption scheme  $(\text{Enc}', \text{Dec}')$  is CCA-secure if for every poly-size  $\mathcal{A}$  there exists a negligible function  $\mu$  such that for every  $\lambda \in \mathbb{N}$ ,  $\mathcal{A}$  wins in the game below with probability at most  $\frac{1}{2} + \mu(\lambda)$ :

1. The challenger chooses a key  $k \leftarrow \mathcal{K}_\lambda$ .
2. The adversary  $\mathcal{A}$  given  $1^\lambda$  can choose a message  $m_i \in \mathcal{M}_\lambda$  and receive  $\text{ct}_i \leftarrow \text{Enc}'(k, m_i)$ .  
Alternatively, he can choose a ciphertext  $\text{ct}_i$  and receive  $m_i = \text{Dec}'(k, \text{ct}_i)$ .  
This step can be repeated polynomially many times, where  $\mathcal{A}$  can query  $\text{Enc}'(k, \cdot)$  and  $\text{Dec}'(k, \cdot)$  polynomially many times.
3. The adversary  $\mathcal{A}$  chooses  $m_0, m_1 \in \mathcal{M}_\lambda$ .
4. The challenger chooses a random bit  $b \leftarrow \{0, 1\}$ , generates  $\text{ct}_b \leftarrow \text{Enc}'(k, m_b)$ , and sends the ciphertext  $\text{ct}_b$  to the adversary.
5. Again  $\mathcal{A}$  is given oracle access to  $\text{Enc}'(k, \cdot)$  and  $\text{Dec}'(k, \cdot)$ , but he is not allowed to send the query  $\text{ct}_b$  to the decryption oracle (but it can send any other ciphertext, even one obtained by changing a single bit of  $\text{ct}_b$ ).
6. The adversary outputs a bit  $b'$ .

We say that  $\mathcal{A}$  wins if  $b' = b$ .

**Claim 1.** Fix any MAC that is existentially unforgeable against adaptive chosen message attacks, and any CPA-secure encryption scheme  $(\text{Enc}, \text{Dec})$ . Then the authenticated encryption scheme Enc-then-MAC, defined above, is CCA-secure.

*Remark.* Note that if we change the authenticated encryption scheme to be MAC-then-Enc the resulting scheme would not necessarily be CCA-secure.

*Proof.* Fix any MAC that is existentially unforgeable against adaptive chosen message attacks, and any CPA-secure encryption. Suppose there exists a poly-size adversary  $\mathcal{A}$  that wins in the CCA game with non-negligible probability  $\epsilon$ . Denote by  $m_1, \dots, m_\ell$  all the queries that  $\mathcal{A}$  sent to its encryption oracle, and denote by  $(c_1, \tau_1), \dots, (c_\ell, \tau_\ell)$  the answers it obtains. Denote all the queries that  $\mathcal{A}$  sends to the decryption oracle by  $ct'_1, \dots, ct'_\ell$ , and for every  $i \in [\ell]$  parse  $ct'_i = (c'_i, \tau'_i)$ .

For every  $\lambda$  denote by

$$\delta(\lambda) := \Pr[\exists i \in [\ell] : (\text{MAC}(k_{\text{MAC}}, c'_i) = \tau'_i) \wedge (c'_i \notin \{c_1, \dots, c_\ell\})], \quad (1)$$

where the probability above is over the randomness of the CCA game.

We distinguish between two cases:

- **Case 1:**  $\delta$  is a non-negligible function. In this case we can use  $\mathcal{A}$  to break the security of the MAC scheme.

Specifically, we construct a poly-size adversary  $\mathcal{B}$  and prove that it wins in the MAC security game with probability  $\geq \frac{\delta}{\ell}$ , which is non-negligible.  $\mathcal{B}$  does the following:

1. Sample a random key  $k_{\text{Enc}} \leftarrow \{0, 1\}^\lambda$  for the CPA-secure encryption scheme.
2. Emulate  $\mathcal{A}$ 's oracles, i.e., the encryption and decryption oracles corresponding to the authenticated encryption, using  $k_{\text{Enc}}$  and the oracle to  $\text{MAC}(k_{\text{MAC}}, \cdot)$ .
3. Denote by  $ct'_1, \dots, ct'_\ell$  the oracle calls that  $\mathcal{A}$  makes to its decryption oracle.
4. Choose at random  $i \leftarrow [\ell]$ , and parse  $ct'_i = (c'_i, \tau'_i)$ .
5. Output  $(c'_i, \tau'_i)$

By Equation (1) it holds that with probability  $\geq \frac{\delta}{\ell}$  both  $\tau'_i$  is a valid tag of  $c'_i$  and  $\mathcal{B}$  didn't send  $c'_i$  as an oracle query.

- **Case 2:**  $\delta$  is a negligible function. In this case, we argue that the decryption algorithm is useless, in the sense that we can simulate the decryption algorithm without knowing the secret key. Formally, we show how to convert  $\mathcal{A}$  into a poly-size adversary  $\mathcal{B}$  that breaks the CPA security of the underlying encryption scheme  $(\text{Enc}, \text{Dec})$ .

$\mathcal{B}$  emulates  $\mathcal{A}$ , as follows:

1. a set  $S = \emptyset$ .
2. Sample  $k_{\text{MAC}} \leftarrow \{0, 1\}^\lambda$ .

3. Every time that  $\mathcal{A}$  queries its encryption oracle with a message  $m \in \mathcal{M}_\lambda$ ,  $\mathcal{B}$  does the following:
  - (a) Send  $m$  to the encryption oracle to obtain  $c \leftarrow \text{Enc}(k_{\text{Enc}}, m)$ .
  - (b) Compute  $\tau = \text{MAC}(k_{\text{MAC}}, c)$ .
  - (c) Let  $S = S \cup \{(m, c, \tau)\}$ .
4. Every time that  $\mathcal{A}$  queries its decryption oracle with a ciphertext  $\text{ct} = (c, \tau)$ ,  $\mathcal{B}$  does the following:
  - (a) Check if there exists  $m$  such that  $(m, c, \tau) \in S$ . If so, emulate the decryption oracle by outputting  $m$ .
  - (b) Otherwise, output  $\perp$ .
5. Once  $\mathcal{A}$  chooses the two challenge messages  $m_0, m_1 \in \mathcal{M}_\lambda$ ,  $\mathcal{B}$  forwards these challenge messages  $m_0, m_1$  to the challenger.
6. Upon receiving  $c_b \leftarrow \text{Enc}(k_{\text{Enc}}, m_b)$  for a random  $b \leftarrow \{0, 1\}$ ,  $\mathcal{B}$  proceeds to simulate the encryption and decryption oracles of  $\mathcal{A}$  as above, and finally it outputs the guess  $b'$  that  $\mathcal{A}$  outputs.

The fact that we are in Case 2 implies that there exists a negligible function  $\delta$  such that for every  $\lambda \in \mathbb{N}$ ,

$$\Pr[b' = b] \geq \epsilon(\lambda) - \delta(\lambda),$$

contradicting the assumption that  $(\text{Enc}, \text{Dec})$  is CPA-secure.  $\square$

## References