EE 720: An Introduction to Number Theory and Cryptography (Spring 2020)

Lecture 12 — March 2, 2020

Instructor: Saravanan Vijayakumaran *Scribe: Saravanan Vijayakumaran*

1 Lecture Plan

- Authenticated Encryption

2 Authenticated Encryption

- In many applications where secrecy is needed it turns out that integrity is also essential. It is a best practice to ensure *both secrecy and integrity* of the messages.
- Let $\Pi = (\texttt{Gen}, \texttt{Enc}, \texttt{Dec})$ be private-key encryption scheme.
  - We require Π to be CCA-secure.
  - We want the scheme to satisfy existential unforgeability under an adaptive chosen-message attack
- But Π does not have the syntax of a message authentication code. So we introduce a definition of integrity for this case.
- **The unforgeable encryption experiment** $\texttt{Enc-Forge}_{\mathcal{A},\Pi}(n)$:
  1. Run $\texttt{Gen}(1^n)$ to obtain a key k.
  2. The adversary $\mathcal{A}$ is given input 1^n and access to an encryption oracle $\texttt{Enc}_k(\cdot)$. The adversary outputs a ciphertext c.
  3. Let $m := Dec_k(c)$, and let $\mathcal{Q}$ denote the set of all queries that $\mathcal{A}$ asked its encryption oracle. The output of the experiment is 1 if and only if (1) $m \neq \perp$ and (2) $m \notin \mathcal{Q}$.

Definition. *A private-key encryption scheme* Π *is* ***unforgeable*** *if for all PPT adversaries* $\mathcal{A}$*, there is a negligible function* $\texttt{negl}$ *such that:*

$$\Pr\left[\texttt{Enc-Forge}_{\mathcal{A},\Pi}(n) = 1\right] \leq \texttt{negl}(n).$$

Definition. *A private-key encryption scheme* Π *is an* ***authenticated encryption scheme*** *if it is CCA-secure and unforgeable.*

2.1 How to Construct AE Schemes?

- One should be careful while combining a secure encryption scheme and secure MAC to construct a secure authenticated encryption scheme.

- Let $\Pi_E = (\mathtt{Enc}, \mathtt{Dec})$ be a CPA-secure encryption scheme and let $\Pi_M = (\mathtt{Mac}, \mathtt{Vrfy})$ denote a MAC, where key generation in both schemes involves choosing a uniform n-bit key.[1]

- Assume independent keys k_E and k_M for Π_E and Π_M, respectively. There are three natural approaches to combining encryption and authentication.

  1. *Encrypt-and-authenticate:* Given a plaintext message m, the sender transmits ciphertext $\langle c, t\rangle$ where:
  $$c \leftarrow \mathtt{Enc}_{k_E}(m) \text{ and } t \leftarrow \mathtt{Mac}_{k_M}(m).$$
  The receiver decrypts c to recover m; assuming no error occurred, it then verifies the tag t.

  2. *Authenticate-then-encrypt:* A MAC tag t is first computed, and then the message and tag are encrypted together. Given a message m, the sender transmits the ciphertext c computed as:
  $$t \leftarrow \mathtt{Mac}_{k_M}(m) \text{ and } c \leftarrow \mathtt{Enc}_{k_E}(m\|t).$$
  The receiver decrypts c to recover $m\|t$; assuming no error occurred, it then verifies the tag t.

  3. *Encrypt-then-authenticate:* Given a plaintext message m, the message is first encrypted and then a MAC tag is computed over the result. The ciphertext is the pair $\langle c, t\rangle$ where:
  $$c \leftarrow \mathtt{Enc}_{k_E}(m) \text{ and } t \leftarrow \mathtt{Mac}_{k_M}(c).$$
  The receiver first verifies the tag t; assuming no error occurred it decrypts c to recover m.

- The first two approaches are not secure.

  - In encrypt-and-authenticate, if a *deterministic* MAC is used then the resulting scheme is not CPA-secure.
  - In authenticate-then-encrypt, the ciphertext c is not authenticated and is vulnerable to padding oracle attacks. There are now two sources of decryption failure: the padding may be incorrect or the MAC tag may not verify. Even if one error is returned for both failures, timing attacks can be used to figure out which failure occurred.

- The encrypt-then-authenticate scheme is an authenticated encryption scheme if the MAC is *strongly secure*.

- A secure MAC guarantees that any PPT adversary will not be able to forge a valid tag t for a message $m \notin \mathcal{Q}$ (the set of messages queried). The definition does not say anything about the situation when the adversary can generate a different valid tag $t' \neq t$ for some message in $\mathcal{Q}$. A secure MAC allows such an adversary to exist. But a strongly secure MAC excludes such adversaries.

- Consider the experiment $\mathtt{Mac\text{-}sforge}_{\mathcal{A},\Pi}(n)$:

  1. A key k is generated by running $\mathtt{Gen}(1^n)$.

[1] Note that we do not assume Π_E is CCA-secure. We have not seen a construction of such a scheme. The point of introducing MACs was to construct CCA-secure schemes.

2. The adversary $\mathcal{A}$ is given input 1^n and oracle access to $\mathtt{Mac}_k(\cdot)$. The adversary eventually outputs (m, t). Let $\mathcal{Q}$ denote the set of all pairs (m', t') of oracle queries and their associated responses.
3. $\mathcal{A}$ succeeds if and only if (1) $\mathtt{Vrfy}_k(m, t) = 1$ and (2) $(m, t) \notin \mathcal{Q}$. If $\mathcal{A}$ succeeds, the output of the experiment is 1. Otherwise, the output is 0.

Definition. *A message authentication code* $\Pi = (\mathtt{Gen}, \mathtt{Mac}, \mathtt{Vrfy})$ *is* ***strongly secure****, or* ***a strong MAC****, if for all PPT adversaries* $\mathcal{A}$*, there is a negligible function* $\mathtt{negl}$ *such that:*

$$\Pr\left[\mathit{Mac\text{-}sforge}_{\mathcal{A},\Pi}(n) = 1\right] \leq \mathtt{negl}(n).$$

Proposition. *Let* $\Pi = (\mathtt{Gen}, \mathtt{Mac}, \mathtt{Vrfy})$ *be a secure MAC that uses canonical verification. Then* Π *is a strong MAC.*

- **Encrypt-then-authenticate construction:** Let $\Pi_E = (\mathtt{Enc}, \mathtt{Dec})$ be a CPA-secure encryption scheme and let $\Pi_M = (\mathtt{Mac}, \mathtt{Vrfy})$ denote a MAC, where key generation in both schemes involves choosing a uniform n-bit key. Define a private-key encryption schem $(\mathtt{Gen}', \mathtt{Enc}', \mathtt{Dec}')$ as follows:
  - $\mathtt{Gen}'$: on input 1^n, choose independent, uniform $k_E, k_M \in \{0,1\}^n$ and output the key (k_E, k_M).
  - $\mathtt{Enc}'$: on input a key (k_E, k_M) and a plaintext message m, compute $c \leftarrow \mathtt{Enc}_{k_E}(m)$ and $t \leftarrow \mathtt{Mac}_{k_M}(c)$. Output the ciphertext $\langle c, t \rangle$.
  - $\mathtt{Enc}'$: on input a key (k_E, k_M) and a ciphertext $\langle c, t \rangle$, first check $\mathtt{Vrfy}_{k_M}(c, t) = 1$. If yes, then output $\mathtt{Dec}_{k_E}(c)$; if no, then output $\perp$.

Theorem. *Let* Π_E *be a CPA-secure private-key encryption scheme, and let* Π_M *be a strongly secure MAC. Then the above construction is an authenticated encryption scheme.*

- To motivate why we get an authenticated encryption from the above construction, we will say that a ciphertext $\langle c, t \rangle$ is *valid* if t is a valid MAC tag on c. Strong security of the MAC means that the adversary will not be able to generate any valid ciphertext that it did not receive from the encryption oracle. This implies that the scheme is unforgeable.
- CCA-security follows because the strong MAC on the ciphertext makes the decryption oracle useless. For every ciphertext $\langle c, t \rangle$ that the adversary submits to the decryption oracle, it either already knows the decryption or it gets an error due to authentication failure. But it already knows the tag failure will occur if it modified the ciphertext. So no new information is gained by the adversary.

3 References and Additional Reading

- Sections 4.2, 4.5 from Katz/Lindell