EE 720: Introduction to Number Theory and Cryptography (Autumn 2023)
Instructor: Saravanan Vijayakumaran
Indian Institute of Technology Bombay

Endsemester Exam: 40 points — Date: November 21, 2023

1. [5 points] Let $F : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}^n$ be a length-preserving pseudorandom function. Consider the private-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ for messages of length n as follows:

   - Gen: On input 1^n, choose k uniformly from $\{0,1\}^n$.
   - Enc: Given $k \in \{0,1\}^n$ and message $m \in \{0,1\}^n$, choose uniform $r \in \{0,1\}^n$ and output the ciphertext
   $$c := \langle r, F_k(r) \oplus m \rangle.$$
   - Dec: Given $k \in \{0,1\}^n$ and ciphertext $c = \langle r, s \rangle$, output the plaintext message
   $$m := F_k(r) \oplus s.$$

   Show that Π is **not** CCA-secure.

   Note 1: The *CCA indistinguishability experiment* $\mathsf{PrivK}^{\mathsf{cca}}_{\mathcal{A},\Pi}(n)$ where $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is described below.

   1. A key k is generated by running $\mathsf{Gen}(1^n)$.
   2. The adversary $\mathcal{A}$ is given 1^n and oracle access to $\mathsf{Enc}_k(\cdot)$ and $\mathsf{Dec}_k(\cdot)$. It outputs a pair of messages $m_0, m_1 \in \mathcal{M}$ with $|m_0| = |m_1|$, where $\mathcal{M}$ is the message space.
   3. A uniform bit $b \in \{0,1\}$ is chosen. Ciphertext $c \leftarrow \mathsf{Enc}_k(m_b)$ is computed and given to $\mathcal{A}$. c is called the **challenge ciphertext**.
   4. The adversary $\mathcal{A}$ continues to have oracle access to $\mathsf{Enc}_k(\cdot)$ and $\mathsf{Dec}_k(\cdot)$, but is not allowed to query the latter on the challenge ciphertext itself. Eventually, $\mathcal{A}$ outputs a bit b'.
   5. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise. If output is 1, we say that $\mathcal{A}$ succeeds.

   Note 2: A private-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is **CCA-secure**, if for all PPT adversaries $\mathcal{A}$ there is a negligible function negl such that, for all n,
   $$\Pr\left[\mathsf{PrivK}^{\mathsf{cca}}_{\mathcal{A},\Pi}(n) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(n).$$

2. [5 points] A private-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is **EAV-secure**, if for all PPT adversaries $\mathcal{A}$ there is a negligible function negl such that, for all n,
   $$\Pr\left[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n) = 1\right] \leq \frac{1}{2} + \mathsf{negl}(n).$$

   Let $\mathsf{out}_{\mathcal{A}}\left(\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n,b)\right)$ denote the output b' of $\mathcal{A}$ when m_b is encrypted. Suppose that a private-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is EAV-secure.

   Prove that for all PPT adversaries $\mathcal{A}$ there is a negligible function negl such that, for all n,
   $$\left|\Pr\left[\mathsf{out}_{\mathcal{A}}\left(\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n,0)\right) = 1\right] - \Pr\left[\mathsf{out}_{\mathcal{A}}\left(\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n,1)\right) = 1\right]\right| \leq \mathsf{negl}(n).$$

   Note: The $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n)$ experiment is obtained by removing the encryption and decryption oracle access to $\mathcal{A}$ in the $\mathsf{PrivK}^{\mathsf{cca}}_{\mathcal{A},\Pi}(n)$ experiment (described in the note after question 1).

3. [5 points] Let $F : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}^n$ be a length-preserving pseudorandom function. Using F, construct a CCA-secure private-key encryption scheme for messages of length n. State the results which lead to the CCA security of your scheme. You don't have to prove these stated results.

   Note: CBC-MAC is a secure deterministic MAC for fixed-length messages that uses canonical verification. It consists of a triple of algorithms $(\mathsf{Gen}, \mathsf{Mac}, \mathsf{Vrfy})$ that operate as follows.

   - Let $m \in \{0,1\}^{dn}$ be a message for a fixed $d > 0$.
   - Gen: On input 1^n, choose k uniformly from $\{0,1\}^n$. This key is assumed to be available to both sender and receiver.
   - Mac:
     1. Parse the message m in to d blocks $m_1, \ldots, m_d$ of length n bits each.
     2. Set $t_0 = 0^n$. For $i = 1, \ldots, d$, set
     $$t_i = F_k(t_{i-1} \oplus m_i).$$
     3. Output t_d as the tag.
   - Vrfy: For a message-tag pair (m, t) output 0, if the message is not of length dn. Otherwise, output 1 if and only if $t = \mathsf{Mac}_k(m)$.

4. Alice claims to be good at predicting the future. Consider the following scenarios.

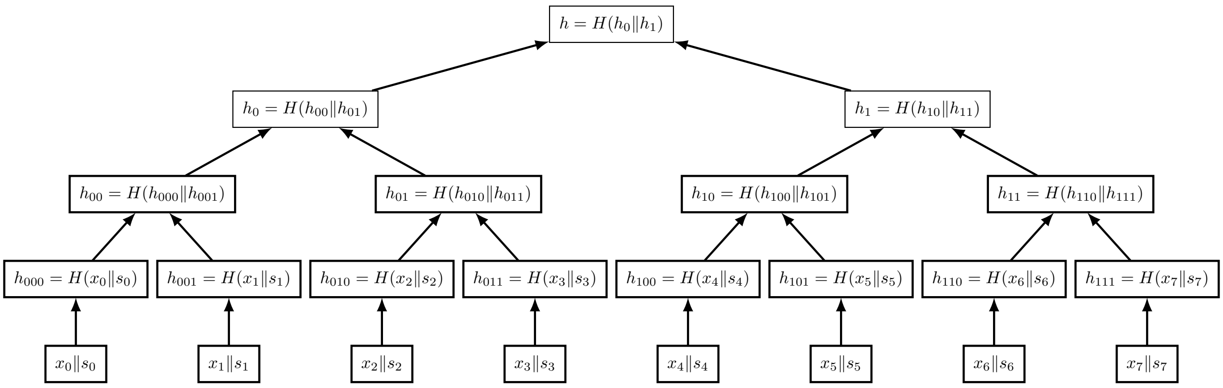

(a) [2½ points] She claims that she can predict the top 8 eight teams out of the 10 teams and their ordering after the league stage of the ICC World Cup **before** the tournament begins. Alice wants to keep her prediction secret (maybe she is married to a player and does not want to cause a controversy). But she wants to prove her claim. Alice does the following.

- Let x_0 correspond to the name of the top team, x_1 correspond to the name of the team with the second highest points, and so on.
- Alice constructs a Merkle tree having 8 leaves as shown in the figure. The function H is assumed to be a cryptographic hash function like SHA256. The values $s_0, s_1, \dots, s_7$ are random n-bit strings which act like salt values, where $n \geq 256$.
- Alice publishes the root h of the Merkle tree on social media like Twitter **before** the tournament begins. Without the salt values, someone who knows h can try out all $10 \times 9 \times \cdots \times 3$ possibilities (≈ 2 million) for the ordering of teams to figure out Alice's prediction.
- At some later point, Alice **only** wants to prove that leaf x_i had a particular value. She does not want to reveal the other leaves or their locations in the list.

What is the **minimum amount of information** Alice needs to reveal to prove that a leaf x_i was present at position i in the list?

- The name of a team can be counted as one unit of information.
- A hash value can be counted as one unit of information.
- A salt value can also be counted as one unit of information.
- The position i is part of the statement to be proved, i.e. x_i is the leaf at position i, and need **not** be included in the information revealed by Alice.

(b) [2½ points] Now suppose Alice's son attends a coaching class for 12th standard students which has 1000 students. She claims she can predict the ordering of the board exam marks of the students **before** the results are declared. Once again, she wants to keep her prediction secret to avoid affecting the performance of the students. For example, if a student is ranked last by Alice, he may get discouraged and stop working hard for the board exams.

How can Alice modify the strategy described in the previous part to publish a root hash h on social media and later prove that that leaf x_i had a particular value?

In this case, what is the **minimum amount of information** Alice needs to reveal to prove that a leaf x_i was present at position i in the list? Assume that salt values are appended to student names to create the leaves.

5. [5 points] Let a, b be integers not both zero. Let c also be an integer. Prove that the equation $ax + by = c$ has a solution (x, y) in $\mathbb{Z}^2$ if and only if $\gcd(a, b)$ divides c.

6. [5 points] Use the Chinese remainder theorem to find all solutions of the following equation in $\mathbb{Z}_{187}$.

$$x^2 + 3x + 2 = 0 \bmod 187.$$

7. [5 points] Suppose a message $m \in \mathbb{Z}_{713}^*$ is encrypted using plain RSA two times. The first time the encryption exponent $e_1 = 3$ is used and the second time the encryption exponent $e_2 = 257$ is used. The ciphertexts in the two cases were $c_1 = 711$ and $c_2 = 313$ respectively. Find the message m. Show your steps. **Note:** $713 = 23 \times 31$.

8. [5 points] Find the four square roots of 187 in $\mathbb{Z}_{713}^*$. Show your steps. **Note:** $713 = 23 \times 31$.