EE 720: An Introduction to Number Theory and Cryptography (Spring 2020)
Instructor: Saravanan Vijayakumaran
Indian Institute of Technology Bombay

Midsem Exam : 25 points | February 27, 2020

1. (5 points) Let $f : \{0,1\}^n \to \{0,1\}^n$ be any function. The Feistel transform of f is the function $FSTL_f : \{0,1\}^{2n} \to \{0,1\}^{2n}$ defined by
$$FSTL_f(L\|R) = R\|f(R) \oplus L$$
where L and R both belong to $\{0,1\}^n$, $\oplus$ denotes the bitwise XOR operator, and $\|$ denotes the string concatenation operator.

   Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ be a length-preserving pseudorandom function. Define $F' : \{0,1\}^n \times \{0,1\}^n \times \{0,1\}^{2n} \to \{0,1\}^{2n}$ as
$$F'(k_1, k_2, L\|R) = FSTL_{F_{k_2}}\left(FSTL_{F_{k_1}}(L\|R)\right)$$
   If k_1, k_2 are chosen independently and uniformly from $\{0,1\}^n$, prove that $F'(k_1, k_2, \cdot) : \{0,1\}^{2n} \to \{0,1\}^{2n}$ is **not** a pseudorandom permutation. Note that the distinguisher knows the structure of F' but does not have access to the keys k_1, k_2.

   Note: $F'(k_1, k_2, \cdot)$ is a pseudorandom permutation if for every PPT distinguisher D, there is a negligible function $\texttt{negl}$ such that:
$$\left|\Pr\left[D^{F'(k_1,k_2,\cdot)}(1^n) = 1\right] - \Pr\left[D^{g(\cdot)}(1^n) = 1\right]\right| \leq \texttt{negl}(n),$$
   where the first probability is taken over uniform choice of $k_1, k_2 \in \{0,1\}^n$ and the randomness of D, and the second probability is taken over uniform choice of $g \in \texttt{Perm}_{2n}$ and the randomness of D. The set $\texttt{Perm}_{2n}$ is the set of all bijections with domain equal to $\{0,1\}^{2n}$ and range equal to $\{0,1\}^{2n}$. By $D^{F'(k_1,k_2,\cdot)}(1^n)$ and $D^{g(\cdot)}(1^n)$, we mean distinguishers D who have oracle access to $F'(k_1, k_2, \cdot)$ and g respectively.

2. (a) (2½ points) Let $\Pi = (\texttt{Gen}, \texttt{Enc}, \texttt{Dec})$ be a private-key encryption scheme with message space $\mathcal{M} = \{0,1\}^n$ where n is the security parameter. Let $\mathcal{C}$ be the ciphertext space of Π. Define $\Pi' = (\texttt{Gen}', \texttt{Enc}', \texttt{Dec}')$ to be another private-key encryption scheme with message space $\mathcal{M}' = \{0,1\}^{2n}$ as follows:
   - $\texttt{Gen}'$ is the same algorithm as $\texttt{Gen}$. Both algorithms output a key k from a keyspace $\mathcal{K}$.
   - For key k and message $m = m_1\|m_2 \in \{0,1\}^{2n}$ where $m_1 \in \{0,1\}^n$ and $m_2 \in \{0,1\}^n$, the ciphertext output by $\texttt{Enc}'$ is $(c_1, c_2) \in \mathcal{C} \times \mathcal{C}$ where $c_1 \leftarrow \texttt{Enc}_k(m_1)$ and $c_2 \leftarrow \texttt{Enc}_k(m_2)$.
   - For key k and ciphertext $(c_1, c_2) \in \mathcal{C} \times \mathcal{C}$, the message output by $\texttt{Dec}'$ is $m_1\|m_2$ where $m_1 = \texttt{Dec}_k(c_1)$ and $m_2 = \texttt{Dec}_k(c_2)$.

   Prove that Π' is **not** CCA-secure, even if Π is CCA-secure.

   (b) (2½ points) Let $\Pi_1 = (\texttt{Gen}_1, \texttt{Enc}_1, \texttt{Dec}_1)$ and Let $\Pi_2 = (\texttt{Gen}_2, \texttt{Enc}_2, \texttt{Dec}_2)$ be two private-key encryption schemes with the same message space $\mathcal{M}$ and ciphertext space $\mathcal{C}$. Define $\Pi' = (\texttt{Gen}', \texttt{Enc}', \texttt{Dec}')$ to be another private-key encryption scheme with message space $\mathcal{M}$ and ciphertext space $\mathcal{C} \times \mathcal{C}$ as follows:
   - The key generated by $\texttt{Gen}'$ is (k_1, k_2) where $k_1 \leftarrow \texttt{Gen}_1(1^n)$ and $k_2 \leftarrow \texttt{Gen}_2(1^n)$.
   - For key (k_1, k_2) and message $m \in \mathcal{M}$, the ciphertext output by $\texttt{Enc}'$ is $(c_1, c_2) \in \mathcal{C} \times \mathcal{C}$ where $c_1 \leftarrow \texttt{Enc}_{1,k_1}(m)$ and $c_2 \leftarrow \texttt{Enc}_{2,k_2}(m)$.
   - For key (k_1, k_2) and ciphertext $(c_1, c_2) \in \mathcal{C} \times \mathcal{C}$, the decryption algorithm $\texttt{Dec}'$ first computes $m_1 = \texttt{Dec}_{1,k_1}(c_1)$ and $m_2 = \texttt{Dec}_{2,k_2}(c_2)$. If $m_1 \neq m_2$, $\texttt{Dec}'$ outputs $\perp$ to indicate decryption failure. If $m_1 = m_2$, $\texttt{Dec}'$ outputs m_1.

   Prove that Π' is **not** CCA-secure, even if Π_1 and Π_2 are CCA-secure.

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

1. A key k is generated by running $\texttt{Gen}(1^n)$.
2. The adversary $\mathcal{A}$ is given 1^n and oracle access to $\texttt{Enc}_k(\cdot)$ and $\texttt{Dec}_k(\cdot)$. It outputs a pair of messages $m_0, m_1 \in \mathcal{M}$ with $|m_0| = |m_1|$.
3. A uniform bit $b \in \{0,1\}$ is chosen. Ciphertext $c \leftarrow \texttt{Enc}_k(m_b)$ is computed and given to $\mathcal{A}$. c is called the *challenge ciphertext*.
4. The adversary $\mathcal{A}$ continues to have oracle access to $\texttt{Enc}_k(\cdot)$ and $\texttt{Dec}_k(\cdot)$, 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 $\Pi = (\texttt{Gen}, \texttt{Enc}, \texttt{Dec})$ has **indistinguishable encryptions under a chosen-ciphertext attack**, or is **CCA-secure**, if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there is a negligible function $\texttt{negl}$ such that, for all n,

$$\Pr\left[\texttt{PrivK}^{\texttt{cca}}_{\mathcal{A},\Pi}(n) = 1\right] \leq \frac{1}{2} + \texttt{negl}(n).$$

3. (a) (2½ points) Show that the CBC block cipher mode encryption scheme is not CCA-secure.

   (b) (2½ points) Show that the CTR block cipher mode encryption scheme is not CCA-secure.

   Note 1: Cipher Block Chaining (CBC) mode works as follows:

   - Let $m = m_1, m_2, \ldots, m_l$ where $m_i \in \{0,1\}^n$.
   - Let F be a length-preserving pseudorandom permutation with block length n.
   - A uniform *initialization vector (IV)* of length n is first chosen.
   - Set $c_0 = IV$. For $i = 1, \ldots, l$, set $c_i := F_k(c_{i-1} \oplus m_i)$. The ciphertext is $(c_0, c_1, \ldots, c_l)$.
   - For $i = 1, 2, \ldots, l$, $m_i := F_k^{-1}(c_i) \oplus c_{i-1}$.

   Note 2: Counter (CTR) mode works as follows:

   - Let $m = m_1, m_2, \ldots, m_l$ where $m_i \in \{0,1\}^n$.
   - Let F be a length-preserving pseudorandom function with block length n.
   - A uniform value $\texttt{ctr}$ of length n is first chosen.
   - Set $c_0 = \texttt{ctr}$. For $i = 1, \ldots, l$, set $c_i := F_k(\texttt{ctr} + i) \oplus m_i$. The ciphertext is $(c_0, c_1, \ldots, c_l)$.
   - For $i = 1, 2, \ldots, l$, $m_i := F_k(\texttt{ctr} + i) \oplus c_i$.

4. Recall that the PKCS #5 padding scheme is used to pad a message x having length some integral number of bytes into a *encoded data* m having length jL bytes where L is the block length in bytes. The number of bytes which are appended to x to get m is b where $1 \leq b \leq L$. Each of these padding bytes is equal to the byte representation of the integer b. Assume that $L < 256$.

   Suppose the encoded data m has length $2L$ bytes, i.e. $m = (m_1, m_2)$ where $|m_i| = L$ bytes for $i = 1, 2$. Recall that the encoded data m is obtained by padding the message x. Let F be a length-preserving pseudorandom permutation where $F_k : \{0,1\}^n \mapsto \{0,1\}^n$ where $n = 8L$. (***Note:*** *$8L$ bits = L bytes*)

   Now suppose the encoded data is encrypted using CBC mode as described below.

   - The ciphertext corresponding to $m = (m_1, m_2)$ is given by $c = (c_0, c_1, c_2)$ where
     - c_0 is uniformly chosen from $\{0,1\}^n$.
     - $c_i = F_k(m_i \oplus c_{i-1})$ for $i = 1, 2$.

   Suppose an adversary has access to a padding oracle. On input some ciphertext block c' of the form (c'_0, c'_1, c'_2) or (c'_0, c'_1), the padding oracle only returns a message from the set $\{\texttt{ok}, \texttt{padding\_error}\}$. The $\texttt{ok}$ is returned when there is no padding error in the encoded data m' obtained from c'. If there is a padding error, then $\texttt{padding\_error}$ is returned.

   (a) (2 points) Describe a procedure by which the adversary can recover the **length** b of the padding in the encoded data m. Be specific about the inputs sent to the padding oracle and the decisions made by your procedure on receiving the oracle's responses.

   (b) (3 points) Describe a procedure by which the adversary can recover **all the message bytes** in the encoded data m. The adversary is allowed to send ciphertext blocks of the form $c' = (c'_0, c'_1)$ or of the form $c' = (c'_0, c'_1, c'_2)$.

5. (5 points) Recall the construction of a CBC-MAC using a length-preserving pseudorandom function F with key/input/output length equal to n bits. Let $m \in \{0,1\}^{dn}$ be a message for a fixed integer $d > 1$.

   - $\texttt{Gen}$: Choose key k uniformly from $\{0,1\}^n$.
   - $\texttt{Mac}$: Parse the message m into d blocks $m_1, \ldots, m_d$ of length n bits each.
     Set $t_0 = 0^n$. For $i = 1, \ldots, d$, set $t_i = F_k(t_{i-1} \oplus m_i)$.
     Output t_d as the tag.
   - $\texttt{Vrfy}$: For a message-tag pair (m, t) output 0, if the message is not of length dn.
     Otherwise, output 1 if and only if $t = \texttt{Mac}_k(m)$.

Consider the following variation of the CBC-MAC. Prove that this variation $(\mathtt{Gen}, \mathtt{Mac}', \mathtt{Vrfy}')$ is an **insecure** MAC.

- $\mathtt{Mac}'$: Parse the message m in to d blocks $m_1, \ldots, m_d$ of length n bits each.
  Choose an initialization vector IV uniformly from $\{0,1\}^n$.
  Set $t_0 = IV$. For $i = 1, \ldots, d$, set $t_i = F_k(t_{i-1} \oplus m_i)$.
  Output (IV, t_d) as the tag.
- $\mathtt{Vrfy}'$: For a message-tag pair $(m, (IV, t))$ output 0, if the message is not of length dn.
  Parse the message m into d blocks $m_1, \ldots, m_d$ of length n bits each.
  Set $t_0 = IV$. For $i = 1, \ldots, d$, set $t_i = F_k(t_{i-1} \oplus m_i)$.
  Output 1 if and only if $t = t_d$.

Note: A message authentication code $\Pi = (\mathtt{Gen}, \mathtt{Mac}, \mathtt{Vrfy})$ is existentially unforgeable under an adaptive chosen-message attack, or just **secure**, if for all PPT adversaries $\mathcal{A}$, there is a negligible function $\mathtt{negl}$ such that:

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

The message authentication experiment $\mathtt{Mac\text{-}forge}_{\mathcal{A},\Pi}(n)$ is defined as follows:

1. A key k is generated by running $\mathtt{Gen}(1^n)$.
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 queries that $\mathcal{A}$ asked its oracle.
3. $\mathcal{A}$ succeeds if and only if (1) $\mathtt{Vrfy}_k(m, t) = 1$ and (2) $m \notin \mathcal{Q}$. If $\mathcal{A}$ succeeds, the output of the experiment is 1. Otherwise, the output is 0.