TD 5: PRFs (corrected version)

Exercise 1. *CTR Security*

Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ be a PRF. To encrypt a message $M \in \{0,1\}^{d \cdot n}$, CTR proceeds as follows:

- Write $M = M_0 \| M_1 \| \ldots \| M_{d-1}$ with each $M_i \in \{0,1\}^n$.
- Sample IV uniformly in $\{0,1\}^n$.
- Return $IV \| C_0 \| C_1 \| \ldots \| C_{d-1}$ with $C_i = M_i \oplus F(k, IV + i \bmod 2^n)$ for all i.

The goal of this exercise is to prove the security of the CTR encryption mode against chosen plaintext attacks, when the PRF F is secure.

1. Recall the definition of security of an encryption scheme against chosen plaintext attacks.

☞ Let $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ be an encryption scheme. We consider the following experiments Exp_b for $b \in \{0,1\}$:

- Challenger samples $k \leftarrow \mathsf{KeyGen}$,
- Adversary makes q encryption queries on messages $(M_{i,0}, M_{i,1})$,
- Challenger sends back $Enc(k, M_{i,b})$ for each i,
- Adversary returns $b' \in \{0,1\}$.

We define the advantage of the adversary $\mathcal{A}$ against the encryption scheme as

$$\mathsf{Adv}^{\mathsf{CPA}}(\mathcal{A}) = \left|\Pr(\mathcal{A} \xrightarrow{\mathsf{Exp}_1} 1) - \Pr(\mathcal{A} \xrightarrow{\mathsf{Exp}_0} 1)\right|.$$

Then, the encryption scheme is said to be secure against chosen plaintext attacks if no probabilistic polynomial-time adversary has a non-negligible advantage with respect to n.

(Note in particular that since $\mathcal{A}$ runs in polynomial time, q must be polynomial in n.)

Remark: in another equivalent definition, there is only one experiment in which the challenger starts by choosing the bit b uniformly at random, and the advantage is defined as $\mathsf{Adv}^{\mathsf{CPA}}(\mathcal{A}) = |\Pr(\mathcal{A} \to 1 \mid b = 0) - \Pr(\mathcal{A} \to 1 \mid b = 1)|$.

2. Assume an attacker makes Q encryption queries. Let $IV_1, \ldots, IV_Q$ be the corresponding IV's. Let `Twice` denote the event "there exist $i, j \leq Q$ and $k_i, k_j < d$ such that $IV_i + k_i = IV_j + k_j \bmod 2^n$ and $i \neq j$." Show that the probability of `Twice` is bounded from above by $Q^2 d / 2^{n-1}$.

☞ *Remark: the probability of* `Twice` *is obviously 1 if it is not required that i and j be distinct. Besides, considering the case $i = j$ is not interesting for our purpose.*

For $i, j \leq Q$, let $\mathtt{Twice}_{i,j}$ be the event "$\exists k_i, k_j < d : \mathrm{IV}_i + k_i = \mathrm{IV}_j + k_j \pmod{2^n}$", which is equivalent to "$\exists k, |k| < d$ and $\mathrm{IV}_i - \mathrm{IV}_j = k \pmod{2^n}$". As the IVs are chosen uniformly and independently, $\mathrm{IV}_i - \mathrm{IV}_j$ is uniform modulo 2^n and $\Pr(\mathtt{Twice}_{i,j}) \leq 2^{-n}(2d-1)$. (The inequality is strict when $2d - 1 > 2^n$, in which case $\Pr(\mathtt{Twice}_{i,j}) = 1$.) Then,

$$\Pr(\mathtt{Twice}) \leq \sum_{1 \leq i \neq j \leq Q} \Pr(\mathtt{Twice}_{i,j}) = Q(Q-1)2^{-n}(2d-1) \leq 2^{1-n} Q^2 d.$$

3. Assume the PRF F is replaced by a uniformly chosen function $f : \{0,1\}^n \to \{0,1\}^n$. Give an upper bound on the distinguishing advantage of an adversary $\mathcal{A}$ against this idealized version of CTR, as a function of d, n and the number of encryption queries Q.

☞ We write $M^{i,\beta} = M_0^{i,\beta} \| \ldots \| M_{d-1}^{i,\beta}$ with $1 \leq i \leq Q$ and $\beta \in \{0,1\}$ the encryption queries of the adversary $\mathcal{A}$ and $C^i = \mathrm{IV}_i \| C_0^i \| \ldots \| C_{d-1}^i$ with $1 \leq i \leq Q$ the replies. Given the value of $b \in \{0,1\}$ chosen by the challenger, we know that $C_j^i = M_j^{i,b} \oplus f(\mathrm{IV}_i + j \pmod{2^n})$ for all $1 \leq i \leq Q$ and $0 \leq j < d$.

If `Twice` does not occur, then all the $\mathrm{IV}_i + j \pmod{2^n}$ for $1 \leq i \leq Q$ and $0 \leq j < d$ are pairwise distinct. Then the values of f at these points are independent and uniformly distributed, since $f : \{0,1\}^n \to \{0,1\}^n$ is chosen uniformly at random. Therefore, all the C_j^i are also independent and uniformly distributed regardless of the value of b, so that $\Pr(\neg\mathtt{Twice} \wedge \mathcal{A} \to 1 \mid b = 0) = \Pr(\neg\mathtt{Twice} \wedge \mathcal{A} \to 1 \mid b = 1)$. It follows that

$$\begin{aligned}
\mathsf{Adv}_{\mathcal{U}}^{\mathsf{CPA}}(\mathcal{A}) &= |\Pr(\mathtt{Twice} \wedge \mathcal{A} \to 1 \mid b = 0) - \Pr(\mathtt{Twice} \wedge \mathcal{A} \to 1 \mid b = 1)| \\
&= |\Pr(\mathcal{A} \to 1 \mid b = 0, \mathtt{Twice}) - \Pr(\mathcal{A} \to 1 \mid b = 1, \mathtt{Twice})| \Pr(\mathtt{Twice}) \\
&\leq \Pr(\mathtt{Twice}) \leq 2^{1-n} Q^2 d.
\end{aligned}$$

4. Show that if there exists a probabilistic polynomial-time adversary $\mathcal{A}$ against CTR based on PRF F, then there exists a probabilistic polynomial-time adversary $\mathcal{B}$ against the PRF F. Give a lower bound on the advantage degradation of the reduction.

☞ Assume that $\mathcal{A}$ is a PPT adversary against the encryption scheme with a non-negligible advantage for a chosen plaintext attack. We build an adversary $\mathcal{B}$ against the underlying PRF F as follows:

   1. Choose $b \in \{0,1\}$ uniformly at random.
   2. For each encryption query (M^0, M^1) from $\mathcal{A}$, encrypt M^b using the given scheme, that is,
      (a) Choose $\mathrm{IV} \in \{0,1\}^n$ uniformly at random.
      (b) For $j = 0$ to $d-1$, send a query for $\mathrm{IV} + j$ and with the reply f_j compute $C_j = M_j^b \oplus f_j$.
      (c) Send $\mathrm{IV} \| C_0 \| \dots \| C_{d-1}$ back to $\mathcal{A}$.
   3. When $\mathcal{A}$ finally outputs a bit $b' \in \{0,1\}$, output 1 if $b' = b$ and 0 otherwise.

The advantage of $\mathcal{B}$ against the PRF F is

$$\mathsf{Adv}_F^{\mathsf{PRF}}(\mathcal{B}) = |\Pr(\mathcal{B} \to 1 \mid \mathsf{PRF}) - \Pr(\mathcal{B} \to 1 \mid \ \mathsf{Unif})|$$

where PRF is the experiment in which replies to $\mathcal{B}$ are computed by calling F and Unif is the one in which replies to $\mathcal{B}$ are computed from a uniformly chosen random function f.

Considering the two terms separately gives

$$\begin{aligned}\Pr(\mathcal{B} \to 1 \mid E) &= \frac{1}{2}\left(\Pr(b' = 0 \mid E, b = 0) + \Pr(b' = 1 \mid E, b = 1)\right) \\ &= \frac{1}{2}\left(1 + \Pr(\mathcal{A} \to 1 \mid E, b = 1) - \Pr(\mathcal{A} \to 0 \mid E, b = 0)\right)\end{aligned}$$

where E is either PRF or Unif. Therefore

$$\mathsf{Adv}_F^{\mathsf{PRF}}(\mathcal{B}) \geq \frac{1}{2}\left(\mathsf{Adv}^{\mathsf{CPA}}(\mathcal{A}) - \mathsf{Adv}_{\mathcal{U}}^{\mathsf{CPA}}(\mathcal{A})\right) \geq \frac{1}{2}\mathsf{Adv}^{\mathsf{CPA}}(\mathcal{A}) - 2^{1-n}Q^2 d$$

using the previous question. Thus, if $\mathsf{Adv}^{\mathsf{CPA}}(\mathcal{A})$ is non-negligible then so is $\mathsf{Adv}_F^{\mathsf{PRF}}(\mathcal{B})$, which is then about a half of $\mathsf{Adv}^{\mathsf{CPA}}(\mathcal{A})$.

Exercise 2. *PRF from DDH*

Let $n \in \mathbb{N}$ be a security parameter. Let $\mathbb{G}$ be a cyclic group of prime order $q > 2^n$ which is generated by a public $g \in \mathbb{G}$ and for which DDH is presumably hard.
We want to build a secure Pseudo-Random Function (PRF) under the DDH assumption in $\mathbb{G}$. The following construction was proposed by Naor and Reingold in 1997.
We define the function $F : \mathbb{Z}_q^{n+1} \times \{0,1\}^n \to \mathbb{G}$ as:

$$F(K, x) = g^{a_0 \cdot \prod_{j=1}^{n} a_j^{x_j}},$$

where we parsed $K = (a_0, a_1, \dots, a_n)^\top$ and $x = (x_1, x_2, \dots, x_n)^\top$.
For an index $i \in [1, n]$, we consider an experiment where the adversary is given oracle access to a hybrid function $F^{(i)}(K, \cdot)$ such that

$$\forall x \in \{0,1\}^n, F^{(i)}(K, x) = g^{R^{(i)}(x[1\dots i]) \cdot \prod_{j=i+1}^{n} a_j^{x_j}},$$

where $R^{(i)} : \{0,1\}^i \to \mathbb{Z}_q$ is a uniformly sampled function and $x[1 \dots i]$ denotes the i first bits of x.

1. Prove that in the adversary's view, $F^{(0)}$ behaves exactly as the function F if we define $x[1 \dots 0] = \varepsilon$, the empty string. How does $F^{(n)}$ behave in the adversary's view?

☞ Define $a_0 := R(\varepsilon)$. This value is uniformly sampled over $\mathbb{Z}_q$ since R is uniformly sampled. Then for any key $K \hookleftarrow U(\mathbb{Z}_q^{n+1})$ sampled by the challenger at the beginning, if we define $K' := (a_0, K[1, \dots, n])\top$, then K' is still uniformly sampled and $F^{(0)}(K, \cdot) = F(K', \cdot)$, which does not change the adversary's view.

In the case of $F^{(n)}$, for any $x \in \{0,1\}^n$, $F^{(n)}(K, x) = g^{R(x)}$, which is uniformly distributed over $\mathbb{G}$.

2. Let (g^a, g^b, g^c) be a DDH instance, where $a, b \hookleftarrow U(\mathbb{Z}_q)$ and we have to decide whether $c = ab$ or if $c \hookleftarrow U(\mathbb{Z}_q)$. Describe a probabilistic polynomial-time algorithm that creates Q randomized instances of DDH $\{g^a, g^{b_\ell}, g^{c_\ell}\}_{\ell=1}^{Q}$, where $\{b_\ell\}_{\ell=1}^{Q}$ are uniformly random and independent over $\mathbb{Z}_q$, with the properties that:

- If $c = ab \bmod q$, then $c_\ell = ab_\ell$ for any $\ell \in [1, q]$.
- If $c \neq ab \bmod q$, then $(b_1, c_1, \ldots, b_Q, c_Q)$ follows the uniform distribution over $(\mathbb{Z}_q)^{2Q}$.

☞ Let x_ℓ, y_ℓ be uniform independent variables over $\mathbb{Z}_q$ for $\ell \in \{1, \ldots, Q\}$. Let $b_\ell := bx_\ell + y_\ell$ and $c_\ell := cx_\ell + ay_\ell$.

First, we can compute g^{b_ℓ} and g^{c_ℓ} in polynomial time: we compute $(g^b)^{x_\ell} \cdot g^{y_\ell}$ and $(g^c)^{x_\ell} \cdot (g^a)^{y_\ell}$.

Assume that $c = ab$. Then $c_\ell = abx_\ell + ay_\ell = a(bx_\ell + y_\ell) = ab_\ell$. Moreover b_ℓ is uniformly distributed as y_ℓ is uniformly distributed, thus we get DDH samples.

Otherwise, if $c \neq ab \bmod q$, we see that we map the vector $(x_\ell, y_\ell)^\top$ to $\begin{pmatrix} b & 1 \\ c & a \end{pmatrix} (x_\ell, y_\ell)^\top$. Notice that the matrix is invertible since $c \neq ab \bmod q$. Then the distribution of c_ℓ and b_ℓ is uniform over $\mathbb{Z}_q^2$ and is independent from any of the other DDH samples.

3. For each $i \in [0, n]$, define the experiment Exp_i where $\mathcal{A}$ is given oracle access to $F^{(i)}(K, \cdot)$ for $K \hookleftarrow U(\mathbb{Z}_q^{n+1})$. After at most Q evaluation queries, $\mathcal{A}$ outputs a bit b'. Prove that for each $i \in [0, n-1]$ it holds that Exp_i is computationally indistinguishable from Exp_{i+1} under the DDH assumption.

☞ Assume that there exists some adversary $\mathcal{A}$ that distinguishes between Exp_i and Exp_{i+1} with non-negligible advantage for some $i \in [0, n-1]$. Let us build $\mathcal{B}$ an adversary against the DDH assumption that does the following.

1. On input (g^a, g^b, g^c), adversary $\mathcal{B}$ samples $a_j \hookleftarrow U(\mathbb{Z}_q)$ for $j = i+2$ to n.
2. Adversary $\mathcal{B}$ samples $(g^a, g^{b_\ell}, g^{c_\ell})$ as in the previous question.
3. Adversary $\mathcal{B}$ creates an empty list L and sets $\alpha := 1$.
4. Adversary $\mathcal{B}$ runs $\mathcal{A}$. When $\mathcal{A}$ queries an input x, adversary $\mathcal{B}$ checks its list L.
   - If there exists (g_1, g_2, g_3) such that $(x[1 \ldots i], (g_1, g_2, g_3)) \in L$, recover (g_1, g_2, g_3).
   - Otherwise, set $(g_1, g_2, g_3) := (g^a, g^{b_\alpha}, g^{c_\alpha})$ and add $(x[1 \ldots i], (g_1, g_2, g_3))$ to L and increase α by one.
5. It outputs $g_2^{\prod_{j=i+2}^n a_j^{x_j}}$ if $x_{i+1} = 0$. Otherwise it outputs $g_3^{\prod_{j=i+2}^n a_j^{x_j}}$.
6. Eventually $\mathcal{A}$ outputs a bit b' that $\mathcal{B}$ outputs too.

We claim that in the case where $c = ab$, the view of $\mathcal{A}$ is the same as if it were given access to $F^{(i)}(K, \cdot)$ and in the case where $c \neq ab$ the view of $\mathcal{A}$ is the same as if it were given access to $F^{(i+1)}(K, \cdot)$ (for uniform K).

Note that we can choose the values of K and R, as long as they are distributed accordingly to Exp_i.

We prove the first part of our claim. Assume that $c = ab$. Since a is uniformly sampled, we can set $K = (a_0, \ldots a_n)^\top$ and $a_{i+1} = a$: the key is still uniformly sampled over $\mathbb{Z}_q^{n+1}$.

Moreover, we can set $b_\alpha = R(x[1 \ldots i])$ where $(x[1 \ldots i], g^a, g^{b_\alpha}, g^{c_\alpha}) \in L$ (by construction, such a α is unique). In that case, since $g^{c_\alpha} = g^{a_{i+1} \cdot b_\alpha}$, it holds that the output of the query is $g^{R(x[1\ldots i]) \cdot \prod_{j=i+1}^n a_j^{x_j}}$, which is exactly $F^{(i)}(K, \cdot)$.
When $c \neq ab$, define R the following way: $R(x[1 \ldots i]0) := b_\alpha$ and $R(x[1 \ldots i]1) := c_\alpha$. This definition of R is valid as every b_α and c_α are uniform and independent. Then, it holds that the output of the query is $g^{R(x[1\ldots i+1]) \cdot \prod_{j=i+2}^n a_j^{x_j}}$ and this gives oracle access to $F^{(i+1)}(K, \cdot)$ to $\mathcal{A}$, with $K := (a_0, \ldots, a_n)^\top$ and the first i a_k are unused. As such, the advantage of $\mathcal{B}$ is:

$$\begin{aligned}
\mathsf{Adv}(\mathcal{B}) &= |\Pr(\mathcal{B} \text{ outputs } 1 | DDH) - \Pr(\mathcal{B} \text{ outputs } 1 | Unif)| \\
&\geq |\Pr(\mathcal{B} \to 1 | c = ab) - \Pr(\mathcal{B} \to 1 | c \neq ab)| - 1/q \\
&\geq \mathsf{Adv}(\mathcal{A}) - 1/q.
\end{aligned}$$

Then $\mathcal{B}$ has non-negligible advantage.

4. Conclude by giving an upper bound on the advatange of a PRF distinguisher as a function of the maximal advantage of a DDH distinguisher.

☞ Assuming the advantage of a DDH distinguisher is at most ε, the advantage of a PRF distinguisher is bounded from above by

$$\mathsf{Adv}(PRF) \leq n \cdot (\varepsilon + 1/q).$$

Remark: Contrary to the GGM construction, the advantage loss does not depend on Q. This is a consequence of the random self-reducibility.