

TD2: Pseudorandom Generators (corrected version)


Exercise 1.*Bit-flip of a PRG*

Let G a pseudo-random generator (PRG) of input range $\{0, 1\}^s$ and output range $\{0, 1\}^n$. We define \bar{G} as follows:

$$\forall x \in \{0, 1\}^s, \bar{G}(x) := 1^n \oplus G(x),$$

where \oplus denotes the XOR operation. This corresponds to flipping every bit of the output of G .

1. Prove that \bar{G} is secure if and only if G is secure.

 Assume that G is secure. We will prove that \bar{G} is secure. Assume by contradiction that there exists an adversary \mathcal{A} that distinguishes between $\bar{G}(U(\{0, 1\}^s))$ and $U(\{0, 1\}^n)$ with non-negligible advantage. We build \mathcal{A}' a distinguisher between $G(U(\{0, 1\}^s))$ and $U(\{0, 1\}^n)$ the following way: on input a sample y , \mathcal{A}' calls \mathcal{A} on the sample $1^n \oplus y$. It outputs the same value.

Notice the following: if y is uniformly distributed, then so is $1^n \oplus y$. If y follows the distribution $G(U(\{0, 1\}^s))$, then $1^n \oplus y$ follows the distribution $\bar{G}(U(\{0, 1\}^s))$. Then \mathcal{A}' 's view is exactly as intended. It guesses from which distribution is sampled $1^n \oplus y$ with non-negligible advantage, and the advantage of \mathcal{A}' is equal to the advantage of \mathcal{A} , which contradicts the assumption that G is secure.

Finally, we notice that the flipped version of \bar{G} is G , and the previous proof also shows that \bar{G} secure implies G secure.

Exercise 2.*Variable-length OTP is not secure*


A variable length one-time pad is a cipher (E, D) , where the keys are bit strings of some fixed length L , while messages and ciphertexts are variable length bit strings, of length at most L . Thus, the cipher (E, D) is defined over $(\mathcal{K}, \mathcal{M}, \mathcal{C})$, where

$$\mathcal{K} := \{0, 1\}^L \text{ and } \mathcal{M} := \mathcal{C} = \{0, 1\}^{\leq L}$$

for some parameter L . Here, $\{0, 1\}^{\leq L}$ denotes the set of all bit strings of length at most L (including the empty string). For a key $k \in \{0, 1\}^L$ and a message $m \in \{0, 1\}^{\leq L}$ of length ℓ , the encryption function is defined as follows:

$$E(k, m) := k[0 \dots \ell - 1] \oplus m$$


1. Provide a counter-example showing that the variable length OTP is not secure for perfect secrecy.

 Suppose that the message distribution contains two messages m_0, m_1 of distinct length, i.e. $|m_0| \neq |m_1|$ in its support. Then given a ciphertext c with $|c| = |m_0|$, we have $\Pr[c = E(k, m_1)] = 0$ while $\Pr_{m \leftarrow \mathcal{M}}[m = m_1] \neq 0$. Hence, the scheme is not perfectly secure.

Exercise 3.

Let $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a function, with $m > n$.

1. Recall the definition of a PRG from the lecture.


 $G : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a PRG if there exists no ppt $\mathcal{A} : \{0, 1\}^m \rightarrow \{0, 1\}$ that distinguish with non-negligible probability between $\mathcal{U}(\{0, 1\}^m)$ and $G(\mathcal{U}(\{0, 1\}^n))$.

Let $\text{Enc} : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^m$ defined by $\text{Enc}(k, m) = G(k) \oplus m$.

2. Give the associated decryption algorithm.

 $\text{Enc} = \text{Dec}$

3. Recall the smCPA security notion from the lecture.

 Two experiments, Exp_b for $b \in \{0, 1\}$ are defined as follows:

1. The challenger \mathcal{C} chooses k uniformly.
2. The adversary \mathcal{A} chooses m_0, m_1 distinct of identical bitlength.
3. The challenger \mathcal{C} returns $\text{Enc}(k, m_b)$.
4. The adversary \mathcal{A} outputs a guess b' .

This is summed up in the following sketch:

\mathcal{C}	\mathcal{A}
$k \leftarrow \mathcal{U}(K)$	
Send $\text{Enc}(k, m_b)$	Choose and send $(m_0, m_1) \in P' := \{(m, m') \in P^2, m \neq m' \wedge m = m' \}$
	Output $b' \in \{0, 1\}$.

The advantage of \mathcal{A} is defined as $\text{Adv}(\mathcal{A}) := |\Pr(\mathcal{A} \xrightarrow{\text{Exp}_0} 1) - \Pr(\mathcal{A} \xrightarrow{\text{Exp}_1} 1)|$. Then (Enc, Dec) is said smCPA-secure if no efficient adversary has non-negligible advantage.

Let $m_1, m_2 \in \{0, 1\}^m$ be arbitrary messages.

- What is the statistical distance between the distributions $\mathcal{U}_1 = m_1 \oplus \mathcal{U}(\{0, 1\}^m)$ and $\mathcal{U}_2 = m_2 \oplus \mathcal{U}(\{0, 1\}^m)$?

☞ They are the same distributions, so 0.

We proved in class that $G \text{ PRG} \Rightarrow (\text{Enc}, \text{Dec}) \text{ smCPA-secure}$. We are going to prove $(\text{Enc}, \text{Dec}) \text{ not smCPA-secure} \Rightarrow G \text{ not PRG}$.

- Let \mathcal{A} be an distinguisher between two games G_0 and G_1 . We say that \mathcal{A} wins if it output 0 (resp 1) during the game G_0 (resp G_1). Show that

$$\text{Adv}_{\mathcal{A}}(G_0, G_1) = 2 \cdot \left| \Pr_{b \sim \mathcal{U}(\{0, 1\})} (\mathcal{A} \text{ wins in } G_b) - \frac{1}{2} \right|$$

☞

$$\Pr_{b \sim \mathcal{U}(\{0, 1\})} (\mathcal{A} \text{ wins in } G_b) = \frac{1}{2} \cdot \Pr_{G_0}(\mathcal{A} \rightarrow 0) + \frac{1}{2} \cdot \Pr_{G_1}(\mathcal{A} \rightarrow 1) = \frac{1}{2} \left(\Pr_{G_0}(\mathcal{A} \rightarrow 0) + 1 - \Pr_{G_1}(\mathcal{A} \rightarrow 0) \right)$$

Hence the result.

- Assume that \mathcal{A} is an adversary with non-negligible advantage ε against the smCPA-security of (Enc, Dec) . Construct an explicit distinguisher between $\mathcal{U}(\{0, 1\}^m)$ and $G(\mathcal{U}(\{0, 1\}^n))$ and compute its advantage.

☞ We define the \mathcal{A}' to be the following:

- Get k from the distribution $G = G(\mathcal{U}(\{0, 1\}^n))$ or $\mathcal{U}(\{0, 1\}^m)$.
- Get m_1, m_2 from \mathcal{A} .
- Sample b from $\mathcal{U}(\{0, 1\})$.
- Send $k \oplus m_b$ to \mathcal{A} and get the output b' .
- If $b = b'$, output "G" else output "U".

The advantage of \mathcal{A}' is $|\Pr_{k \sim G}(\mathcal{A}' \rightarrow G) - \Pr_{k \sim \mathcal{U}}(\mathcal{A}' \rightarrow G)|$.

Assume $k \sim \mathcal{U}$ and define Y_0 the game played when $b = 0$ and Y_1 the game played when $b = 1$. Since $k \sim \mathcal{U}$, we have that $m_b \oplus k$ is independent from m_b , hence $\Pr_{m_0, k}(\mathcal{A}(m_0 \oplus k) \rightarrow 1) = \Pr_{m_1, k}(\mathcal{A}(m_1 \oplus k) \rightarrow 1)$ and hence the advantage of \mathcal{A} between Y_0 and Y_1 is 0. By the previous question we have $\Pr_{b \sim \mathcal{U}(\{0, 1\}), k}(\mathcal{A} \text{ wins when given } m_b \oplus k) = 1/2$.

Assume $k \sim G$ and define Y'_0 the game played when $b = 0$ and Y'_1 the game played when $b = 1$. We have $\Pr_{k \sim G}(\mathcal{A}' \rightarrow G) = \Pr_b(\mathcal{A} \text{ wins } Y'_b)$.

Finally, $\text{Adv}_{\mathcal{A}'} = |\Pr_{k \sim G}(\mathcal{A}' \rightarrow G) - \Pr_{k \sim \mathcal{U}}(\mathcal{A}' \rightarrow G)| = |\Pr_b(\mathcal{A} \text{ wins } Y'_b) - 1/2| = \varepsilon/2$.

Exercise 4.

Let (Enc, Dec) be an encryption scheme over $K \times P \times \{0, 1\}^n$.

- In this question, we assume that (Enc, Dec) is smCPA-secure. Prove that there exists a smCPA-secure encryption scheme $(\text{Enc}', \text{Dec}')$ such that $G : k \mapsto \text{Enc}'(k, 0)$ is not a secure PRG. *Hint: try to concatenate constant bits to every ciphertext.*

☞ Define $\text{Enc}' : (k, m) \mapsto 1^\ell || \text{Enc}(k, m)$. The decryption algorithm Dec' ignores the first ℓ bits and calls Dec on the remaining ones. We have two things to prove:

- The pair $(\text{Enc}', \text{Dec}')$ is a smCPA-secure encryption scheme.
- $G : k \mapsto 1^\ell || \text{Enc}(k, 0)$ is not a secure PRG.

We start with the first claim. If we assume by contradiction that there exists an efficient adversary \mathcal{A} that breaks the smCPA-security of $(\text{Enc}', \text{Dec}')$, we build \mathcal{A}' against the smCPA-security of (Enc, Dec) the following way. It starts by calling \mathcal{A} . When \mathcal{A} outputs two messages m_0, m_1 , \mathcal{A}' outputs the same messages to the challenger. When the challenger outputs a ciphertext c , \mathcal{A}' sends to \mathcal{A} the ciphertext $1^\ell || c$. When \mathcal{A} outputs a bit b' , \mathcal{A}' outputs the same. This is summed up in the following sketch:

\mathcal{C}	\mathcal{A}'	\mathcal{A}
$k \leftarrow U(K)$	Call \mathcal{A}	
	Send the same messages (m_0, m_1)	Choose and send $(m_0, m_1) \in P'$
Send $c := \text{Enc}(k, m_b)$	Compute and send to \mathcal{A} : $c' := 1^\ell c$	
	Output b'	Output b'

In these games, the view of \mathcal{A} is the same as in the previous question. This means that it behaves the same way as in the Exp_b games for the encryption scheme $(\text{Enc}', \text{Dec})$. By definition of the advantage, $\text{Adv}(\mathcal{A}') = \text{Adv}(\mathcal{A})$. Thus, this breaks the security of (Enc, Dec) .

We move on to prove the second claim by exhibiting an efficient distinguisher \mathcal{B} . It does the following: upon receiving a sample from either $G(U(K))$ or the uniform distribution, it outputs 1 if the first ℓ bits are 1 and 0 otherwise. Its advantage is $1 - \frac{1}{2^\ell}$. It is non-negligible as soon as $\ell \geq 1$.