

Homework 1: Symmetric Cryptography

Due: September 12, 2024 at 11:59pm (Submit on Gradescope)

Instructor: David Wu

Instructions. You **must** typeset your solution in LaTeX using the provided template:

<https://www.cs.utexas.edu/~dwu4/courses/fa24/static/homework.tex>

You must submit your problem set via [Gradescope](#) (accessible through [Canvas](#)).

Collaboration Policy. You may discuss your general *high-level* strategy with other students, but you may not share any written documents or code. You should not search online for solutions to these problems. If you do consult external sources, you must cite them in your submission. You must include the names of all of your collaborators with your submission. Refer to the [official course policies](#) for the full details.

Problem 1: Perfect Secrecy [10 points]. Consider the one-time pad encryption scheme on n -bit messages. Namely, the size of the key, message, and ciphertext are all n -bits long. The encryption of $m \in \{0, 1\}^n$ under $k \in \{0, 1\}^n$ is $c = m \oplus k$. The key is normally sampled as $k \xleftarrow{R} \{0, 1\}^n$. Observe however that if $k = 0^n$, then the ciphertext is simply the message (i.e., $c = m$). This seems like an undesirable property, so consider an alternative version of the one-time pad where $k \xleftarrow{R} \{0, 1\}^n \setminus \{0^n\}$ (i.e., the key is no longer allowed to be 0^n). Does this scheme provide perfect secrecy? Briefly explain in **1-2 sentences** (no formal proof needed).

Problem 2: Understanding Advantage [24 points]. In this problem, we will get some practice with the notion of advantage. Consider the following two experiments between a challenger and an adversary \mathcal{A} :

- **Experiment 0:** In this experiment, the challenger samples a bit $\beta \xleftarrow{R} \{0, 1\}$ and gives β to the adversary. The adversary \mathcal{A} outputs a bit $b' \in \{0, 1\}$ which is the output of the experiment.
- **Experiment 1:** In this experiment, the challenger gives $\beta = 0$ to the adversary. The adversary \mathcal{A} outputs a bit $b' \in \{0, 1\}$ which is the output of the experiment.

For an adversary \mathcal{A} and a bit $b \in \{0, 1\}$, we define $W_b = \Pr[\mathcal{A} \text{ outputs } 1 \text{ in Experiment } b]$. We then define the advantage of \mathcal{A} to be $\text{Adv}[\mathcal{A}] = |W_0 - W_1|$. If the advantage is 1, then that means algorithm \mathcal{A} can perfectly distinguish between the two experiments. If the advantage is 0 for all adversaries \mathcal{A} , then the two experiments are identical. If the advantage of every efficient adversary is negligible (with respect to some security parameter), then we say that the two experiments are computationally indistinguishable.

(a) Compute the advantage for each of the following adversaries:

- Algorithm \mathcal{A} always outputs 1.
- Algorithm \mathcal{A} outputs $b' \xleftarrow{R} \{0, 1\}$.
- Algorithm \mathcal{A} outputs β .
- Algorithm \mathcal{A} outputs $1 - \beta$.
- Algorithm \mathcal{A} outputs 1 if $\beta = 1$ and $b' \xleftarrow{R} \{0, 1\}$ if $\beta = 0$.

- (b) What is the maximum possible advantage of any adversary for distinguishing between these two experiments. Give a formal proof of this. **Hint:** One way to model an *arbitrary* adversary is to consider its input/output behavior. Namely, let $p_b := \Pr[\mathcal{A} \text{ outputs } 1 \mid \beta = b]$. Try writing the advantage of \mathcal{A} as a function of p_0 and p_1 . What is the maximum value of this function?

Problem 3: Pseudorandom Generators [25 points]. Let $G: \{0, 1\}^\lambda \rightarrow \{0, 1\}^n$ be a secure PRG. For each of the following functions G' , indicate whether it is a secure PRG or not. If it is secure, give a *formal* proof. If not, describe an explicit attack.

- (a) $G'(s) := G(s) \parallel (G(s) \oplus 1^n)$, where 1^n denotes the all-ones string of length n .
- (b) $G'(s_1 \parallel s_2) := G(s_1) \oplus G(s_2)$.
- (c) $G'(s_1 \parallel s_2) := s_1 \parallel G(s_2)$, where $s_1 \in \{0, 1\}^\lambda$.

Recall that for two bit-strings $s_1, s_2 \in \{0, 1\}^*$, we write $s_1 \parallel s_2$ to denote the *concatenation* of s_1 and s_2 . Please refer to this [handout](#) for examples of how to formally show whether a construction is secure or not.

Problem 4: Encrypting Twice [10 points]. Intuitively, encrypting a message twice should not harm security. It turns out that this is not always true. Let $(\text{Encrypt}, \text{Decrypt})$ be a cipher and define the “encrypt-twice” cipher $(\text{Encrypt}_2, \text{Decrypt}_2)$ where $\text{Encrypt}_2(k, m) := \text{Encrypt}(k, \text{Encrypt}(k, m))$. Give an example of a cipher $(\text{Encrypt}, \text{Decrypt})$ that is semantically secure, but $(\text{Encrypt}_2, \text{Decrypt}_2)$ is not semantically secure. You can state results from lecture *without* proof.

Problem 5: Time Spent [1 point]. How long did you spend on this problem set? This is for calibration purposes, and the response you provide does not affect your score.

Optional Feedback. Please answer the following *optional* questions to help us design future problem sets. You do not need to answer these questions. However, we do encourage you to provide us feedback on how to improve the course experience.

- (a) What was your favorite problem on this problem set? Why?
- (b) What was your least favorite problem on this problem set? Why?
- (c) Do you have any other feedback for this problem set?
- (d) Do you have any other feedback on the course so far?