Foundations of Cryptography, Fall 2025
Problem Set 1
Due Friday, September 19

Total Number of Points: 100

Collaboration Policy: Collaboration is permitted and encouraged in small groups of at most three students. You are free to collaborate in discussing answers, but you must write up solutions on your own, and must specify in your submission the names of any collaborators. Do not copy any text from your collaborators; the writeup must be entirely your work. Do not write down solutions on a board and copy it verbatim; again, the writeup must be entirely your own words and your own work and should demonstrate clear understanding of the solution. **Solutions should be typeset in LaTeX.** You may make use of published material, provided that you clearly acknowledge all sources/tools used. Of course, scavenging for solutions from prior years is forbidden.

On the Use of LLMs: You may use AI however you wish to deepen your understanding of the lecture material. Upload the notes, talk to your AI about them, ask for more explanation or examples; it's all fine. You may not use LLMs in any way to work on your homework. You may not upload assignments, ask for hints, ask how certain concepts from the lectures might be applied to specific homework problems, or upload your assignments to check for correctness or clarity or anything else. You may not include any AI generated content whatsoever in your homework submissions. If it becomes clear that you have used an AI tool when working on your homework (either directly by making edits or to ask for hints/solutions), we may mark your grade down to reflect that.

Problem 0: Probability Theory Handout (10 points)

Before you begin the problem set, please read the handout on probability theory posted on the website. The material covered on the handout will be assumed throughout this course; indicate here once you have read the handout.

Problem 1: Some Probability Review (15 points)

Consider the process of tossing n balls into n bins, where each ball is thrown independently into a uniformly random bin.

1. (5 points) What is the expected number of collisions?[1]
2. (10 points) Prove that for sufficiently large n, the maximum number of balls in a bin is at most $\frac{3\log n}{\log\log n}$ with probability at least $1 - 1/n$.

[1] The number of collisions are counted over distinct but unordered pairs; if exactly three balls are in a bin, then this contributes three collisions to the total count.

Problem 2: Negligible Functions (20 points)

Let $f, g : \mathbb{N} \to \mathbb{R}$ be two negligible functions such that $f(n) \geq g(n)$. Which of the following functions are necessarily negligible?

1. (5 points) $f'(n) := \frac{1}{2^{(\log n)^2}}$.
2. (5 points) $f'(n) := f(n) + g(n)$.
3. (5 points) $f'(n) := \frac{g(n)}{f(n)}$.
4. (5 points) $f'(n) := \sqrt[r]{f(n)}$ for any constant $r > 1$.

Problem 3: One-Way Functions (20 points)

Let $f : \{0,1\}^n \to \{0,1\}^{m(n)}$ and $g : \{0,1\}^n \to \{0,1\}^{p(n)}$ be families of one-way functions (OWFs) where $m(n), p(n) : \mathbb{N} \to \mathbb{N}$ are strictly increasing functions. For each of the following functions, prove or disprove if they are necessarily a OWF. You may assume the existence of length-preserving one-way functions.

1. (5 points) $f'(x) := f(x)_{[1:m/2]}$.
2. (5 points) $f'(x) := f(x_{[1:n/2]})$.
3. (5 points) $f'(x) := f(g(x))$.
4. (5 points) $f'(x) := f(x)||x_{[1:\log n]}$.

Problem 4: Pseudorandom Generators (15 points)

Let $G_1, G_2 : \{0,1\}^n \to \{0,1\}^{m(n)}$ be pseudorandom generators (PRG) where $m(n) > n$. For each of the following functions, prove or disprove if they are necessarily a PRG.

1. (5 points) $G(s) = G_1(s)||G_2(s)$.
2. (5 points) $G(s_1||s_2) = G_1(s_1)||G_2(s_2)$, where $|s_1| = |s_2|$ or $|s_1| = |s_2| + 1$.
3. (5 points) $G(s) = G_1(s||G_1(s))$.

Problem 5: Variations on Shannon Secrecy (20 points)

Ben Bitdiddle is fascinated with Shannon's "perfect secrecy" definition for encryption schemes that he learned in the first lecture of 6.520. However, he is disappointed in some of the shortcomings of this definition of secrecy:

- A *consequence* of this secrecy definition is that any encryption scheme satisfying the definition would require rather long keys, indeed as long as the messages they encrypt. This was proved in Lecture 1.
- The secrecy definition requires *absolute privacy* of the secret key; that is, there are no guarantees whatsoever if an eavesdropper obtains some partial information about the secret key by some means.

Ben's goal is to somehow avoid these shortcomings and your goal is to either help him circumvent each shortcoming or convince him that it cannot be done.

Let us first recall the syntax of an encryption scheme. An encryption scheme with key space $\mathcal{K}$, message space $\mathcal{M}$, and ciphertext space $\mathcal{C}$ consists of two algorithms $(\mathsf{Enc}, \mathsf{Dec})$ described as follows:

- $\mathsf{Enc}(k, m)$ takes as input a key $k \in \mathcal{K}$ and a message $m \in \mathcal{M}$, and produces a ciphertext $c \in \mathcal{C}$.
- $\mathsf{Dec}(k, c)$ takes as input a key $k \in \mathcal{K}$ and a ciphertext $c \in \mathcal{C}$ and produces a message $m \in \mathcal{M}$.

The key for the encryption scheme is chosen from some fixed distribution K over $\mathcal{K}$. You should not assume anything specific about this distribution, besides the fact that it is part of the description of the encryption scheme. In what follows, we assume that the message space $\mathcal{M}$ contains at least two messages.

Shorter Keys? (10 points) We could hope to make do with shorter keys by *relaxing* the secrecy requirement. Here is one possible relaxation: looking at the ciphertext *does* change the *a posteriori* probability of a message, but only a little. An encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is said to have $(1 - \varepsilon)$-secrecy if for any distribution M of messages, for every $m \in \mathcal{M}$ and $c \in \mathcal{C}$,

$$\Pr[M = m | \mathsf{Enc}(K, M) = c] \leq (1 + \varepsilon) \Pr[M = m]$$

As expected, perfect secrecy corresponds to $\varepsilon = 0$.

Prove that relaxing secrecy does not decrease the size of the key space by much. That is, if an encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is perfectly correct and has $(1 - \varepsilon)$-secrecy, show that

$$|\mathcal{K}| \geq \frac{1}{1 + \varepsilon} |\mathcal{M}|.$$

A Leaky Secret Key. (10 points) We now return to the keyed setting, but consider a situation in which the eavesdropper Eve has acquired some information about the secret key k. (This could happen, for instance, if Eve mounts a “side-channel attack” on the device running the encryption/decryption algorithm). In particular, suppose that $\mathcal{K} = \{0, 1\}^n$, and that Eve learns a collection of bits $(k_i)_{i \in S}$ for some *unknown* set $S \subset [n]$ of size $0.9n$.

Design an encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ and specify a distribution K over keyspace $\mathcal{K}$ which achieves *perfect secrecy* for one bit messages. That is, you should describe a scheme and prove that it satisfies the property that for every distribution M on messages, every $m \in \mathcal{M}$, every $c \in \mathcal{C}$, every set $S \subset [n]$ of size $0.9n$, and every choice of $(k_i)_{i \in S}$,

$$\Pr[M = m | \mathsf{Enc}(K, M) = c \text{ and } K_i = k_i \text{ for all } i \in S] = \Pr[M = m].$$