Problem Set 4

Total Number of Points: 40.

Collaboration Policy: Collaboration is allowed 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 them verbatim into LaTeX; again, the writeup must be entirely your own words and your own work and should demonstrate a clear understanding of the solution. Additionally, you may make use of published material, provided that you acknowledge all sources used. Of course, scavenging for solutions from prior years is forbidden.

Problem 1.   Hash It Out. (14 points)

In this problem, we will explore four different definitions of hash functions. We say that a family of functions $\mathcal{H} = \{H_k : \{0,1\}^{2n} \to \{0,1\}^n\}_{k \in \{0,1\}^m}$ is a X hash function, for $\mathsf{X} \in \{$"universal", "weak collision resistant", "universal one-way"[1], "collision resistant"$\}$, if it has the following properties:

- **Efficiency**: Given $k \in \{0,1\}^m$ (where $m = \mathsf{poly}(n)$) and $x \in \{0,1\}^{2n}$, $H_k(x) \in \{0,1\}^n$ can be computed in $\mathsf{poly}(n)$ time.

- **Security**:
  - If $\mathsf{X} =$ "universal": There exists a negligible function negl such that for all $x_1 \neq x_2 \in \{0,1\}^{2n}$,
    $$\mathbf{Pr}\left[k \leftarrow \{0,1\}^m : H_k(x_1) = H_k(x_2) \wedge x_1 \neq x_2\right] \leq \mathsf{negl}(n).$$

  - If $\mathsf{X} =$ "weak collision resistant": For all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that the following holds:
    $$\mathbf{Pr}\left[(k, x_1) \leftarrow \{0,1\}^m \times \{0,1\}^{2n}; \mathcal{A}(k, x_1) \to x_2 : H_k(x_1) = H_k(x_2) \wedge x_1 \neq x_2\right] \leq \mathsf{negl}(n).$$

  - If $\mathsf{X} =$ "universal one-way": For all $x_1 \in \{0,1\}^{2n}$ and for all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that the following holds:
    $$\mathbf{Pr}\left[k \leftarrow \{0,1\}^m; \mathcal{A}(k, x_1) \to x_2 : H_k(x_1) = H_k(x_2) \wedge x_1 \neq x_2\right] \leq \mathsf{negl}(n).$$

  - If $\mathsf{X} =$ "collision resistant": For all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that the following holds:
    $$\mathbf{Pr}\left[k \leftarrow \{0,1\}^m; \mathcal{A}(k) \to (x_1, x_2) : H_k(x_1) = H_k(x_2) \wedge x_1 \neq x_2\right] \leq \mathsf{negl}(n).$$

(a) (2 points) Construct a family of universal hash functions without using any assumptions.
Hint: Consider a random matrix $A \leftarrow \{0,1\}^{n \times 2n}$.

(b) (3 points) Show how to build a universal one-way hash function family from any weak collision resistant hash function family.

(c) (4 points) Show that there exists a hash family that is universal one-way but not collision resistant. (You may assume the existence of universal one-way hash functions.)

---

[1]also called "targeted collision resistant"

(d) (2 points) Show that for all of these definitions, for all $k \in \{0,1\}^m$,

$$\Pr_{x \leftarrow \{0,1\}^{2n}} \left[\left| H_k^{-1}(H_k(x)) \right| \leq 1 \right] \leq \frac{1}{2^n}.$$

(e) (3 points) Suppose $\mathcal{H}$ is a weak collision resistant hash function family. Show that the function $F : \{0,1\}^{m+2n} \to \{0,1\}^{m+n}$ given by $F(k,x) = (k, H_k(x))$ is a one-way function.

It turns out that the existence of one-way functions also implies the existence of weak collision resistant hash functions and universal one-way hash functions, even though it is an open question for collision resistant hash functions! You may assume these facts throughout the problem set.

Problem 2.   Upgrading Lamport Signatures. (8 points)

Recall Lamport's signature scheme from class, based on a OWF $f : \{0,1\}^n \to \{0,1\}^\ell$, that produces an n^2-bit signature for an n-bit message:

$\mathsf{Gen}(1^n)$
___
1 : $x_{1,0}, \ldots, x_{n,0} \leftarrow \{0,1\}^n$

2 : $x_{1,1}, \ldots, x_{n,1} \leftarrow \{0,1\}^n$

3 : $sk := (x_{1,0}, \ldots, x_{n,0}, x_{1,1}, \ldots, x_{n,1})$

4 : $vk := (y_{1,0}, \ldots, y_{n,0}, y_{1,1}, \ldots, y_{n,1})$, where $y_{i,c} = f(x_{i,c})$

5 : **return** (sk, vk)

$\mathsf{Sign}(sk, m \in \{0,1\}^n)$
___
1 : **parse** $sk = (x_{1,0}, \ldots, x_{n,0}, x_{1,1}, \ldots, x_{n,1})$

2 : **return** $\sigma := (x_{1,m_1}, \ldots, x_{n,m_n})$

$\mathsf{Ver}(vk, m \in \{0,1\}^n, \sigma)$
___
1 : **parse** $\sigma := (\sigma_1, \ldots, \sigma_n)$

2 : **if** $\forall i \in [n], f(\sigma_i) \overset{?}{=} y_{i,m_i}$: **return** 1

3 : **else** : **return** 0

In this problem, we will look at a stronger definition of one-time unforgeability known as *one-time strong unforgeability* which states that not only is the adversary unable to produce a signature on a different message, but also that she is unable to produce a *different* signature σ^* on the same message it requested a signature on.

Definition 1 (One-time strong unforgeability)
Let $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Ver})$ *be a digital signature scheme with message space* $\mathcal{M}$ *and key space* $\mathcal{K}$ *with security parameter* n. *This scheme is* <u>*one-time strongly unforgeable*</u> *if for all pair of PPT algorithms* $(\mathcal{A}_1, \mathcal{A}_2)$, *there exists a negligible function* negl *such that for all* n,

$$\Pr \left[\begin{array}{l} (sk, vk) \leftarrow \mathsf{Gen}(1^n); \\ (m, \mathsf{state}) \leftarrow \mathcal{A}_1(vk); \\ \sigma \leftarrow \mathsf{Sign}(sk, m); \\ (m^*, \sigma^*) \leftarrow \mathcal{A}_2(\sigma, \mathsf{state}) \end{array} : \begin{array}{l} (m^*, \sigma^*) \neq (m, \sigma) \wedge \\ \mathsf{Ver}(vk, m^*, \sigma^*) = 1 \end{array} \right] \leq \mathsf{negl}(n).$$

(a) (3 points) Show an attack on the one-time strong unforgeability of Lamport's scheme. That is, assuming the existence of length-preserving one-way functions, construct a one-way function f such that the Lamport signature scheme using f is *not* one-time strongly unforgeable.

(b) (2 points) Prove that if f is an injective one-way function, then Lamport's scheme is one-time strongly unforgeable.

(c) (3 points) Show that if one-way functions exist, then there exists a one-way function such that Lamport's scheme is one-time strongly unforgeable. Note that building injective one-way functions from arbitrary one-way functions is an open question. (You may use statements from class or earlier in this problem set.)

Problem 3. Lossy Encryption. (10 points)

In this problem, we will explore an alternate notion of security for public key encryption schemes called *lossy encryption*. This definition of security is more powerful than IND-CPA security, and allows us to construct other primitives like oblivious transfer and encryption schemes secure against chosen ciphertext attacks.

Lossy encryption schemes have two modes of operation: *real* and *lossy*. In *real* mode, a lossy encryption scheme behaves like a public key encryption scheme. In *lossy* mode, the ciphertexts produced by the encryption algorithm contain *no information* about the message that was encrypted. Formally, a lossy encryption scheme (Gen, Enc, Dec) has the following syntax:

- Gen(1^n, mode): The Gen algorithm takes the security parameter as input (as usual). It also takes as input a mode which can be either real or lossy. In the real mode, it outputs a pair of keys (pk, sk). In the lossy mode, it outputs a lossy public key $\widetilde{pk}$.
- Enc(pk, b): The Enc algorithm takes a public key (either a real or a lossy public key) and a bit b and outputs a ciphertext ct.
- Dec(sk, ct): The Dec algorithm takes as input a secret key (has to be real) and outputs a decrypted bit b.

Furthermore, it has the following properties:

- **Correctness:** The encryption scheme is correct in the real mode. That is, for every bit b,

$$\mathbf{Pr}[(pk, sk) \leftarrow \mathsf{Gen}(1^n, \mathsf{real}) : b = \mathsf{Dec}(sk, \mathsf{Enc}(pk, b))] = 1$$

where the probability is over the randomness of Gen, Enc and Dec algorithms.

- **Key Indistinguishability:** Real public keys are indistinguishable from lossy public keys. That is, for any $n \in \mathbb{N}$,

$$\{pk : (pk, sk) \leftarrow \mathsf{Gen}(1^n, \mathsf{real})\} \approx_c \{\widetilde{pk} : \widetilde{pk} \leftarrow \mathsf{Gen}(1^n, \mathsf{lossy})\},$$

where $\approx_c$ means that the two distributions on the left- and the right-hand sides (the distribution of pk and the distribution of $\widetilde{pk}$, in this case) are computationally indistinguishable.

- **Lossy encryption:** Encryption using the lossy key completely loses information about the message encrypted. That is, output distributions of encryptions of 0 and 1, under lossy keys, are statistically indistinguishable. Formally, for every $n \in \mathbb{N}$, and every $\widetilde{pk}$ in the support of Gen(1^n, lossy),

$$\mathsf{Enc}\left(\widetilde{pk}, 0\right) \approx_s \mathsf{Enc}\left(\widetilde{pk}, 1\right)$$

where the randomness is the coins of the Enc algorithm, and $\approx_s$ means that the two distributions on the left- and right-hand sides are negligibly close (in n) in statistical distance.

(a) (2 points) Show that every lossy encryption scheme also satisfies

$$\left(\widetilde{pk}, \mathsf{Enc}\left(\widetilde{pk}, 0\right)\right) \approx_s \left(\widetilde{pk}, \mathsf{Enc}\left(\widetilde{pk}, 1\right)\right),$$

where $\widetilde{pk} \leftarrow \mathsf{Gen}(1^n, \mathsf{lossy})$, i.e., where $\widetilde{pk}$ is generated randomly according to Gen(1^n, lossy).

(b) (4 points) Show that every lossy encryption scheme is also IND-CPA secure (operating in the real mode).

(c) (4 points) Define a lossy key generation algorithm for the Goldwasser-Micali encryption scheme. Prove the three properties above (correctness, key indistinguishability and lossy encryption) for the Goldwasser-Micali scheme with your lossy key generation algorithm, assuming the Quadratic Residuosity assumption.

Problem 4.   Zero Knowledge Proof for Knight's Move Sudoku (8 points)

Knight's Move Sudoku is a variant on the popular logic puzzle Sudoku. For a Knight's Move Sudoku puzzle for positive integer n, you are given as input an $n^2 \times n^2$ grid of cells, where the cells are partitioned into n^2 disjoint $n \times n$ subgrids. Some of the cells begin filled with integers from 1 to n^2.

A solution is a way to assign an integer from 1 to n^2 to each cell which hasn't already been filled in a way that:

1. Each row contains each of the integers from 1 to n^2 exactly once.

2. Each column contains each of the integers from 1 to n^2 exactly once.

3. Each of the n^2 different $n \times n$ subgrids contains each of the integers from 1 to n^2 exactly once.

4. For all cells that are a "knight's move" away from each other, the entries are distinct. Two cells $(i_1, j_1) \in [n^2] \times [n^2], (i_2, j_2) \in [n^2] \times [n^2]$ are a "knight's move" away if either (a) $|i_1 - i_2| = 1$ and $|j_1 - j_2| = 2$ or (b) $|i_1 - i_2| = 2$ and $|j_1 - j_2| = 1$.
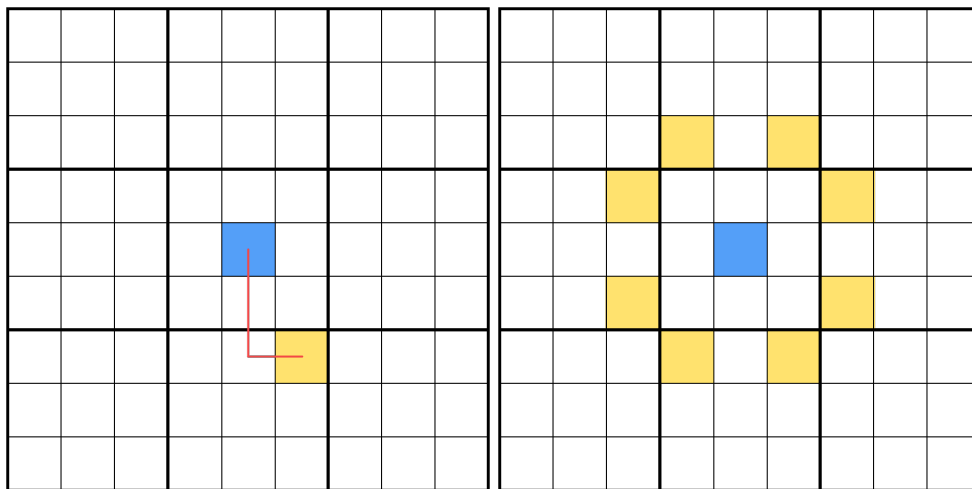

Figure 1: An example demonstrating what a "knight's move" is, for $n = 3$, from https://masteringsudoku.com/chess-sudoku/.

Design a zero-knowledge protocol for Knight's Move Sudoku. Your protocol should be computationally zero knowledge, have perfect completeness, and have soundness error $1 - p(n)$ for a non-negligible function p. You must prove that your construction satisfies all of the above properties. You may assume the existence of one-way functions.