

Quiz 2

Full Marks: 20, Time: 1 hour

Instructor: Sruthi Sekar

Date: 8th Oct 2025

Roll Number: _____ Name: _____

1. Enter both roll number and name in the answer sheet.
2. Answer each question on a new page of the answer booklet.
3. Do not use pencils. Pens only!
4. If you are using the secure construction from any question in another question, feel free to just refer to that question and use it as it is.
5. If you are using a fact proved in class during lectures, please mention the fact clearly and intuitively explain it before using it in your questions.

Problem 1 (2 marks (1+1)). Each part below has a collection of three statements containing two truths and a lie. Can you identify the lie in each part? (No explanations needed!)

- a) Identify the lie amongst the three:
 - i) Every collision-resistant hash function is one-way.
 - ii) Every one-way function is collision-resistant.
 - iii) Every pseudorandom generator is one-way.
- b) Identify the lie amongst the three:
 - i) The hardness of CDH relative to $\mathcal{G}$ implies the hardness of discrete-log relative to $\mathcal{G}$.
 - ii) The hardness of DDH relative to $\mathcal{G}$ implies the hardness of CDH relative to $\mathcal{G}$.
 - iii) The hardness of discrete-log relative to $\mathcal{G}$ implies the hardness of DDH relative to $\mathcal{G}$.

Problem 2 (9 marks (3+3+3)). Are the following functions one-way? If yes, give a formal proof by security reduction, else give a complete attack.

- a) $f(k, m, r) := \text{Enc}(k, m; r) || m$, where (Enc, Dec) is a COA-secure SKE, and r in Enc above refers to the randomness used.
 - b) $f(x) := g(g(x))$, where g is a one-way function.
 - c) $f(x_1 || x_2) := f_1(x_1) \oplus f_2(x_2)$, where f_1 and f_2 are one-way functions and $||$ is concatenation.
-

Problem 3 (3 marks) The discrete log problem is believed to be hard relative to $\mathcal{G}$, where $\mathcal{G}(1^n) \rightarrow (p, g)$ such that p is an n -bit prime and g is the generator of $\mathbb{Z}_p^*$. The function family defined as $f_{p,g}(x) := g^x \bmod p$ is one-way. Let $\text{lsb}(x)$ denote the least-significant bit of x . Show that lsb is not a hard-core predicate of $f_{p,g}$.

Hint: You can assume here that it is easy to decide if a number y is a quadratic residue modulo p or not. No need for an algorithm to show this!

Problem 4 (6 marks (3+3)). Consider *one-wayness* under CPA (OW-CPA), an alternative security notion for PKE defined as follows:

- Eve is given pk , generated as $(pk, sk) \leftarrow \text{Gen}(1^n)$.
- For $m \leftarrow \mathcal{M}_n$, Eve is given $c \leftarrow \text{Enc}(pk, m)$ as the challenge ciphertext.
- Eve outputs m' and breaks if $m' = m$.

A PKE Π is OW-CPA-secure if for all PPT eavesdroppers Eve, the probability with which Eve breaks Π as above is negligible. Answer the following questions about OW-CPA:

- a) I asked ChatGPT to build me a key-exchange protocol secure against an eavesdropper from an OW-CPA PKE, $\Pi_{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$, and it gave me the following construction Π_{KE} that also uses a PRG G to reduce the randomness used by Alice:
 - i) Bob generates $(pk, sk) \leftarrow \text{Gen}(1^n)$ and sends pk to Alice.
 - ii) Alice chooses a random seed $s \in_R \{0, 1\}^n$ and sends $c = \text{Enc}(pk, s)$ to Bob. Alice locally outputs $k_A = G(s)$.
 - iii) Bob computes $\text{Dec}(sk, c) = s$ and outputs $k_B = G(s)$.

Prove that ChatGPT is wrong, i.e., Π_{KE} is insecure against eavesdroppers.

Hint: For your attack, think about using the specific PRG from our class (using a OWP and its hard-core predicate): $G(s) := f_I(x) || hc(x)$, where hc is hard-core predicate of the trapdoor permutation (TDP) family $(\text{Gen}^{\text{TD}}, f, \text{Inv})$. The important thing to use in your attack is the one-wayness of the TDP. Recall, a TDP is just an OWP with the additional guarantee that given the trapdoor td , Inv can efficiently invert!

- b) Can you prove that Π_{KE} is secure against eavesdroppers if Π_{PKE} is IND-CPA secure and G is a PRG. Use a hybrid argument to give this proof. Use can assume the transitivity of computational indistinguishability for this proof.
-

Definitions you need

Discrete Log Assumption: The discrete log problem is hard relative to $\mathcal{G}$ if for all PPT Logan, $\exists$ negligible function negl such that $\forall n$

$$\Pr_{\substack{(\mathbb{G}, q, g) \leftarrow \mathcal{G}(1^n) \\ h \in_R \mathbb{G}}} [\text{Logan}(1^n, (\mathbb{G}, q, g), h) = x : g^x = h] \leq \text{negl}(n)$$

Computational Diffie-Hellman (CDH) Assumption: The CDH problem is hard relative to $\mathcal{G}$ if for all PPT Logan, $\exists$ negligible function negl such that $\forall n$

$$\Pr_{\substack{(\mathbb{G}, q, g) \leftarrow \mathcal{G}(1^n) \\ x, y \in_R \mathbb{Z}_q}} [\text{Logan}(1^n, (\mathbb{G}, q, g), g^x, g^y) = h : h = g^{xy}] \leq \text{negl}(n)$$

Decisional Diffie-Hellman (DDH) Assumption: The DDH problem is hard relative to $\mathcal{G}$ if for all PPT Logan, $\exists$ negligible function negl such that $\forall n$, the following is $\leq \text{negl}(n)$:

$$\left| \Pr_{\substack{(\mathbb{G}, q, g) \leftarrow \mathcal{G}(1^n) \\ x, y \in_R \mathbb{Z}_q}} [\text{Logan}(1^n, (\mathbb{G}, q, g), (g^x, g^y, g^{xy})) = 1] - \Pr_{\substack{(\mathbb{G}, q, g) \leftarrow \mathcal{G}(1^n) \\ x, y, z \in_R \mathbb{Z}_q}} [\text{Logan}(1^n, (\mathbb{G}, q, g), (g^x, g^y, g^z)) = 1] \right|$$

One Way Function (OWFs): $\{f_n : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}\}_{n \in \mathbb{N}}$ is a *one-way function family* if for $\text{Gen}(1^n) \rightarrow f_n$, $f_n(x)$ can be computed efficiently for all x ; and for every PPT algorithm Inv , $\exists$ negl such that $\forall n$

$$\Pr_{x \in_R \{0, 1\}^n} [\text{Inv}(1^n, f_n(x)) \in f_n^{-1}(f_n(x))] \leq \text{negl}(n)$$

Hard-core Predicate (HCP): A function $\text{hc} : \{0, 1\}^n \rightarrow \{0, 1\}$ is a *hard-core predicate* of a function family $\{f_n : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}\}_{n \in \mathbb{N}}$ if hc can be computed in polynomial time, and for every PPT algorithm Pred , $\exists$ negligible function negl such that $\forall n$

$$\Pr_{x \in_R \{0, 1\}^n} [\text{Pred}(1^n, f_n(x)) = \text{hc}(x)] \leq 1/2 + \text{negl}(n)$$

Key Exchange protocol: A two-party key exchange protocol Π is a probabilistic protocol between two parties Alice and Bob, where τ represents the transcript of the protocol, i.e., all the messages exchanged between Alice and Bob during the protocol execution. At the end of the protocol Alice locally outputs $k_A \in \{0, 1\}^n$ and Bob locally outputs $k_B \in \{0, 1\}^n$ such that $\forall n$, $\Pr[k_A = k_B] = 1$. Π is said to be *secure against an eavesdropper* if for all PPT Eve, there exists a negligible function negl such that $\forall n$:

$$\left| \Pr_{(k, \tau) \leftarrow \Pi(1^n)} [\text{Eve}(\tau, k) = 1] - \Pr_{\substack{(k, \tau) \leftarrow \Pi(1^n) \\ r \in_R \{0, 1\}^n}} [\text{Eve}(\tau, r) = 1] \right| \leq \text{negl}(n)$$

Public Key Encryption (PKE): A public key encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ with message space $\mathcal{M}$ satisfies correctness: for every n , for every $(pk, sk) \leftarrow \text{Gen}(1^n)$, and for every $m \in \mathcal{M}$, $\Pr[\text{Dec}(sk, \text{Enc}(pk, m)) = m] = 1$. Recall the following security notion for Π .

- **Chosen Plaintext Attack (CPA):** Define the following game $\text{PubK}_{\text{Eve}, \Pi}^{\text{cpa}}(n)$ between a challenger and the adversary Eve:

1. Challenger generates $(pk, sk) \leftarrow \text{Gen}(1^n)$ and sends pk to Eve. Eve, on input 1^n and pk , outputs a pair of messages m_0, m_1 with $|m_0| = |m_1|$.

2. Challenger picks $b \in_R \{0, 1\}$; sends $c_b \leftarrow \text{Enc}(pk, m_b)$ to Eve.
3. Eve outputs a bit b' .
4. Output 1 if $b' = b$, and 0, otherwise.

Π is said to be *secure against a chosen plaintext attack* if for all PPT adversaries Eve, there exists a negligible function negl such that $\forall n$, $\Pr[\text{PubK}_{\text{Eve}, \Pi}^{\text{cpa}}(n) = 1] \leq 1/2 + \text{negl}(n)$.

Symmetric Key Encryption (SKE): A symmetric-key encryption scheme $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ with message space $\mathcal{M}$ satisfies correctness: for every n , for every $k \leftarrow \text{Gen}(1^n)$, and for every $m \in \mathcal{M}$, $\Pr[\text{Dec}(k, \text{Enc}(k, m)) = m] = 1$. Recall the following security notion for Π .

- **Ciphertext Only Attack (COA):** Define the following game $\text{PrivK}_{\text{Eve}, \Pi}^{\text{coa}}(n)$ between a challenger and the adversary Eve:
 1. Eve, on input 1^n outputs a pair of messages m_0, m_1 with $|m_0| = |m_1|$.
 2. Challenger generates $k \leftarrow \text{Gen}(1^n)$; picks $b \in_R \{0, 1\}$; sends $c_b \leftarrow \text{Enc}(k, m_b)$ to Eve.
 3. Eve on input c_b outputs a bit b' .
 4. Output 1 if $b' = b$, and 0, otherwise.

Π is said to be secure against a ciphertext-only attack if for all PPT adversaries Eve, there exists a negligible function negl such that $\forall n$, $\Pr[\text{PrivK}_{\text{Eve}, \Pi}^{\text{coa}}(n) = 1] \leq 1/2 + \text{negl}(n)$.

Pseudorandom Generators (PRGs): Recall the definition of pseudorandom generators (PRGs). A deterministic polynomial-time algorithm G is a pseudorandom generator (PRG) if it is length-expanding: for any n , and any $k \in \{0, 1\}^n$, $G(k) \in \{0, 1\}^{\ell(n)}$ with $\ell(n) > n$, and satisfies either of the following two equivalent properties:

- **Unpredictability:** G is *next-bit unpredictable* if for every PPT adversary, Eve, there exists a negligible function negl such that $\forall n$

$$\forall i, \Pr[\text{Eve}(G(k)|_{1,2,\dots,i}) = G(k)|_{i+1}] \leq 1/2 + \text{negl}(n),$$

where $G(k)|_{1,2,\dots,i}$ denotes the first i -bits of $G(k)$, and $G(k)|_{i+1}$ denotes its $(i+1)$ -th bit.

- **Pseudorandomness:** G is *pseudorandom* if for every PPT adversary, Eve, there exists a negligible function negl such that $\forall n$

$$\left| \Pr_{k \in_R \{0,1\}^n} [\text{Eve}(G(k)) = 1] - \Pr_{r \in_R \{0,1\}^{\ell(n)}} [\text{Eve}(r) = 1] \right| \leq \text{negl}(n)$$

Collision Resistant Hash Function: A pair of PPT algorithms (Gen, H) with $\text{Gen}(1^n) \rightarrow k$ picks the key k which decides the specific hash function $H^k : \{0, 1\}^* \rightarrow \{0, 1\}^\ell$. (Gen, H) is said to be *collision resistant* if for all PPT adversary Collin, there exists a negligible negl such that $\forall n$,

$$\Pr_{k \leftarrow \text{Gen}(1^n)} [(x, x') \leftarrow \text{Collin}(k) : (H^k(x) = H^k(x')) \wedge (x \neq x')] \leq \text{negl}(n)$$