Midterm I

Name:

SID:

- **You may consult at most *1 double-sided sheet of handwritten notes*. Apart from that, you may not look at books, notes, etc. Calculators, phones, computers, and other electronic devices are NOT permitted for looking up content. However, you may use an electronic device such as a tablet for writing your answers.**
- **DSP Students: If you are allowed $1.5\times$ (resp. $2\times$) the regular exam duration, then you must submit your exam within $130 = 80 * 1.5 + 10$ (resp. $170 = 80 * 2 + 10$) mins.**
- **We will not be answering questions during the exam. If you feel that something is unclear please write a note in your answer.**

1 Multiple Choice (15 points)

In the multiple choice section, no explanations are needed for your answers. Each question is worth 3 points, and there is no penalty for wrong answers. Please mark your answers clearly.

1. Which of the following functions are negligible? (There may be multiple negligible functions.)

   ◯ $2^{-\log_2(n)}$

   ◯ $2^{-(\log_2(n))^3}$

   ◯ $2^{-\sqrt{n}}$

   ◯ $2^{-(n^2)}$

   Solution: The first function is non-negligible because $2^{-\log_2(n)} = \frac{1}{n}$. The rest of the functions are negligible.

2. True or False: If $f(n)$ and $g(n)$ are non-negligible functions, then $h(n) = f(n) \cdot g(n)$ is also non-negligible.

   ◯ True

   ◯ False

   Solution: False. Here is a counterexample:

$$\begin{aligned} \text{Let } f(n) &= \begin{cases} 2^{-n} & , n \text{ is even} \\ 1 & , n \text{ is odd} \end{cases} \\ g(n) &= \begin{cases} 1 & , n \text{ is even} \\ 2^{-n} & , n \text{ is odd} \end{cases} \\ \text{Then } f(n) \cdot g(n) &= 2^{-n} \end{aligned}$$

   $f(n)$ and $g(n)$ are non-negligible, but $f(n) \cdot g(n)$ is negligible.

3. True or False: If an encryption scheme Π is CCA-secure, then it is also CPA-secure.

   ◯ True

   ◯ False

   Solution: True

4. Suppose $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is a CPA-secure encryption scheme that encrypts messages belonging to a field $\mathbb{F}$. Construct a new encryption scheme as follows:

- $\mathsf{Gen}_1(1^n)$: samples $k' \leftarrow \mathsf{Gen}(1^n)$. Then it samples p, a random degree-d polynomial over the field $\mathbb{F}$. The key k for this encryption scheme is the tuple of these values:

$$k = (k', p)$$

(here p refers to the description of the polynomial)

- $\mathsf{Enc}_1(k, m)$ computes and outputs $c = \mathsf{Enc}(k', m) \parallel p(m)$.
- $\mathsf{Dec}_1(k, c)$ just runs $\mathsf{Dec}(k', \cdot)$ on the first part of the ciphertext.

In the CPA security experiment, what is the minimum number of queries to the Enc_1 oracle that are needed to break the CPA security of the scheme $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$. i.e. What is the minimum number of queries needed to figure out b given $\mathsf{Enc}_1(k, m_b)$?

Only count phase-I and phase-II queries; do not count the query used to compute the challenge ciphertext.

Solution: 1 phase-I query is sufficient to break CPA security.

Let us construct an adversary to break CPA security with one phase-I query:

(a) The adversary chooses messages (m_0, m_1) such that $m_0 \neq m_1$, uniformly at random. With high probability, $p(m_0) \neq p(m_1)$.

(b) Then in phase I, they query $\mathsf{Enc}_1(k, m_0)$, so they learn $p(m_0)$.

(c) Then they output challenge messages (m_0, m_1). When they receive the challenge ciphertext, they can check whether it contains $p(m_0)$. If so, they output $b' = 0$. Otherwise, they output $b' = 1$.

If $p(m_0) \neq p(m_1)$, then the adversary is always correct $(b' = b)$.

5. Which of the following modes of encryption do **not** require the PRF/PRP F_k used to be efficiently invertible? There may be multiple such modes.

◯ Electronic Code Book (ECB)

◯ Cipher Block Chaining (CBC)

◯ Output Feedback (OFB)

◯ Counter (CTR)

Solution: OFB, CTR

2 Pseudorandom Functions (15 points)

Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ be a pseudorandom function, and let

$$G(k, (x, y)) = F(k, x) \oplus F(k, y)$$

Prove that G is not a secure pseudorandom function.

The parts below will outline the proof that G is not pseudorandom and ask you to fill in the missing details to complete the proof.

To show that G is not a secure PRF, let us construct a distinguishing algorithm D that can distinguish $G(k, \cdot)$ from a truly random function $R(\cdot)$ (given query access to one of these functions).

1. D makes a single query (x^*, y^*) to the function:

   $(x^*, y^*) = \boxed{}$

   Solution: Choose any arbitrary (x^*, y^*) such that $x^* = y^*$.

   Let z^* be the response obtained.

2. Next, D outputs 1 if

   $\boxed{}$

   and outputs 0 otherwise.

   Solution: D outputs 1 if $z^* = 0^n$.

3. *Pseudorandom case:* In the case where D is querying $G(k, \cdot)$, what is the probability that D outputs 1 (i.e. what is $\Pr[D^{G(k,\cdot)} = 1]$)? Here the probability is over the randomness of D and the randomness of sampling k.

   Explain your reasoning.

   Solution: $\Pr[D^{G(k,\cdot)} = 1] = 1$ because $G(k, (x^*, y^*)) = F(k, x^*) \oplus F(k, x^*) = 0^n$ with certainty.

4. *Truly random case:* Let R be a function sampled uniformly at random from the set of all functions that map $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$.

   In the case where D is querying $R(\cdot)$, what is the probability that D outputs 1 (i.e. what is $\Pr[D^{R(\cdot)} = 1]$)? Here the probability is over the randomness of D and the randomness of sampling R.

Explain your reasoning.

Solution: For any given (x^*, y^*), the value of $R(x^*, y^*)$ is uniformly random over $\{0,1\}^n$, where the randomness is over the choice of R. Therefore, D outputs 1 with probability 2^{-n}.

5. Finish the proof to argue that D breaks PRF security for G.

   Solution: To summarize the previous argument,

$$|\Pr[D^{G(k,\cdot)=1}] - \Pr[D^{R(\cdot)} = 1]| = 1 - 2^{-n}$$

   which is non-negligible. Therefore, G is not a secure PRF.

3 Pseudorandom Functions Again (15 points)

Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ be a pseudorandom function. Prove that the following function H is also a pseudorandom function:

$$H(k, x) = F(k, x) \oplus x$$

Solution:

1. We will prove that if there exists an adversary D_H that distinguishes $H(k, \cdot)$ from $R_1(\cdot)$ with non-negligible advantage, then we can construct an adversary D_F that distinguishes $F(k, \cdot)$ from $R_2(\cdot)$ with the same advantage.

   Description of D_F:

   (a) D_F runs D_H.

   (b) Whenever D_H outputs a query x, D_F forwards the query to its oracle to get a response y. Then it sends to D_H the response $y \oplus x$.

   (c) Finally, D_F outputs whatever D_H outputs.

2. *Pseudorandom case:* If D_F is interacting with an oracle for $F(k, \cdot)$, then it has successfully simulated D_H's interaction with an oracle for $H(k, \cdot)$.

$$\Pr[D_F^{F(k,\cdot)} = 1] = \Pr[D_H^{H(k,\cdot)} = 1]$$

3. *Truly random case:* Let us define

$$R_1(x) = R_2(x) \oplus x$$

   If R_2 is a uniformly random function, then so is R_1. Therefore, if D_F is interacting with an oracle for $R_2(\cdot)$, then it has successfully simulated D_H's interaction with a different random function R_1.

$$\Pr[D_F^{R_2(\cdot)} = 1] = \Pr[D_H^{R_1(\cdot)} = 1]$$

4. In summary:

$$\left| \Pr[D_F^{F(k,\cdot)} = 1] - \Pr[D_F^{R_2(\cdot)} = 1] \right| = \left| \Pr[D_H^{H(k,\cdot)} = 1] - \Pr[D_H^{R_1(\cdot)} = 1] \right|$$

5. Assume toward contradiction that H is not a PRF. Then there exists a D_H such that $\left| \Pr[D_H^{H(k,\cdot)} = 1] - \Pr[D_H^{R_1(\cdot)} = 1] \right|$ is non-negligible. Then we've shown that there exists a D_F such that $\left| \Pr[D_F^{F(k,\cdot)} = 1] - \Pr[D_F^{R_2(\cdot)} = 1] \right|$ is non-negligible. This implies that F is not a PRF, which is contradiction.

6. Therefore our assumption must be false, so in fact, H is a PRF.

Name:

4 CPA-Secure Encryption (20 points)

Let (Gen, Enc, Dec) be a CPA-secure encryption scheme. Below, we will construct another encryption scheme and prove that it is also CPA-secure.

In the encryption scheme below, let the message m belong to $\{0,1\}^n$.

- $\mathsf{Gen}_1(1^n)$: Sample the key as follows: $k \leftarrow \mathsf{Gen}(1^n)$.
- $\mathsf{Enc}_1(k, m)$: Sample $r \leftarrow \{0,1\}^n$ uniformly at random. Then compute $c_0 := \mathsf{Enc}(k, r)$ and $c_1 := r \oplus m$. Output the ciphertext $c = (c_0, c_1)$.
- $\mathsf{Dec}_1(k, (c_0, c_1))$: **Unspecified**

1. Fill in the decryption algorithm so that every ciphertext is decrypted correctly.

   $\mathsf{Dec}_1(k, (c_0, c_1))$:

   Solution: $\mathsf{Dec}_1(k, (c_0, c_1))$: Compute $r' := \mathsf{Dec}(k, c_0)$ and then compute $m' := r' \oplus c_1$. Output m'.

2. Prove that $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ satisfies CPA security.

Solution:

Note: In the solution below, we've included a lot more detail than students would be expected to give on an exam; we believe this makes it easier to learn from the solution. For ease of reading, we've marked in gray the sections that can be skipped if you are just skimming the solution.

1. Let's define two hybrids that are identical, except in eq. (1) and eq. (2) below.
   - $\underline{H_0}$ is the CPA security game for $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$. Also, let $\mathcal{A}$ be the adversary in this game.

     Here's the hybrid in more detail. (This level of detail is optional, but it can be very helpful to the reader of your proofs).

     (a) **Setup:** The challenger samples $k \leftarrow \mathsf{Gen}_1(1^n)$.

(b) **Phase I queries:** $\mathcal{A}$ sends the challenger a message, and the challenger responds with

$$c = \big(\mathsf{Enc}(k, r), (r \oplus m)\big)$$

where $r \leftarrow \{0,1\}^n$. $\mathcal{A}$ can repeat this step many times.

(c) **Challenge:** $\mathcal{A}$ outputs two messages m_0, m_1. The challenger samples a bit $b \leftarrow \{0,1\}$, and sends $\mathcal{A}$ the encryption c^* of m_b:

$$c^* = \mathsf{Enc}(k, r), (r \oplus m_b) \tag{1}$$

where r is sampled uniformly at random.

(d) **Phase II queries:** Work the same as phase I queries.

(e) **Output:** $\mathcal{A}$ outputs a bit b'. The output of the hybrid is 1 if $b = b'$ and 0 otherwise.

- $\underline{H_1}$ is the same as H_0 except the challenge ciphertext c^* is $(\mathsf{Enc}(k, 0^n), (r \oplus m_b))$.

  Here's the hybrid in more detail, with any change from Hybrid 0 underlined:

  (a) **Setup:** The challenger samples $k \leftarrow \mathsf{Gen}_1(1^n)$.

  (b) **Phase I queries:** $\mathcal{A}$ sends the challenger a message, and the challenger responds with

  $$c = \big(\mathsf{Enc}(k, r), (r \oplus m)\big)$$

  where $r \leftarrow \{0,1\}^n$. $\mathcal{A}$ can repeat this step many times.

  (c) **Challenge:** $\mathcal{A}$ outputs two messages m_0, m_1. The challenger samples a bit $b \leftarrow \{0,1\}$, and sends $\mathcal{A}$ the encryption c^* of m_b:

  $$c^* = \mathsf{Enc}(k, \underline{0^n}), (r \oplus m_b) \tag{2}$$

  where r is sampled uniformly at random.

  (d) **Phase II queries:** Work the same as phase I queries.

  (e) **Output:** $\mathcal{A}$ outputs a bit b'. The output of the hybrid is 1 if $b = b'$ and 0 otherwise.

2. **Claim 4.1** *If* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is CPA secure, then for any adversary* $\mathcal{A}$, $\big|\Pr[H_0 \to 1] - \Pr[H_1 \to 1]\big|$ *is negligible.*

   Proof:

   (a) Assume toward contradiction that for some adversary $\mathcal{A}$, $\big|\Pr[H_0 \to 1] - \Pr[H_1 \to 1]\big|$ is non-negligible. Then we will construct an adversary $\mathcal{B}$ that breaks the CPA security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$.

   (b) **Notation:** We'll use (M_0, M_1, B, C^*) to denote some of the variables in the CPA security game in which $\mathcal{B}$ is playing so that they don't get mixed up with (m_0, m_1, b, c^*) from the hybrids above.

   (c) $\mathcal{B}$ is designed to simulate one of the hybrids, H_B, where $B \in \{0,1\}$ is chosen by the CPA challenger.

   At the beginning of the CPA security game, the challenger samples $k \leftarrow \mathsf{Gen}(1^n)$, which serves to simulate step a of the hybrids.

   Description of $\mathcal{B}$:

i. In phase I of the CPA game, $\mathcal{B}$ will simulate step b of the hybrids. This entails running $\mathcal{A}$, and when $\mathcal{A}$ outputs an encryption query m, $\mathcal{B}$ will sample $r \leftarrow \{0,1\}^n$ and send $\mathcal{A}$ the response: $c = (\mathsf{Enc}(k, r), (r \oplus m))$. This requires $\mathcal{B}$ to make a query to $\mathsf{Enc}(k, \cdot)$.

ii. In the challenge phase:

A. $\mathcal{B}$ samples an $r \leftarrow \{0,1\}^n$ and outputs two challenge messages, $M_0 = r, M_1 = 0^n$, and receives in response either $C^* = \mathsf{Enc}(k, r)$ (when $B = 0$) or $C^* = \mathsf{Enc}(k, 0^n)$ (when $B = 1$).

B. Next, C^* allows $\mathcal{B}$ to simulate step c of the hybrids. When $\mathcal{A}$ outputs its messages (m_0, m_1) in step c, $\mathcal{B}$ will sample a bit $b \leftarrow \{0,1\}$ and respond to $\mathcal{A}$ with

$$C^*, (r \oplus m_b)$$

Note that if $C^* = \mathsf{Enc}(k, r)$, then $\mathcal{B}$ has simulated step c of H_0, but if $C^* = \mathsf{Enc}(k, 0^n)$, then $\mathcal{B}$ has simulated step c of H_1.

iii. In phase II, $\mathcal{B}$ will simulate steps d and e of the hybrids. The output of the hybrid is a bit, which $\mathcal{B}$ outputs as well.

(d) As we argued above, $\mathcal{B}$ simulates H_B. Therefore,

$$\begin{aligned}
\Pr[\mathcal{B} \to 1 | B = 0] &= \Pr[H_0 \to 1] \\
\Pr[\mathcal{B} \to 1 | B = 1] &= \Pr[H_1 \to 1] \\
\Big| \Pr[\mathcal{B} \to 1 | B = 0] - \Pr[\mathcal{B} \to 1 | B = 1] \Big| &= \Big| \Pr[H_0 \to 1] - \Pr[H_1 \to 1] \Big| \\
&= \textbf{non-negligible}(n)
\end{aligned}$$

This means that $\mathcal{B}$ breaks the CPA security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$.

(e) However, the problem states that $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is CPA-secure, so we've arrived at a contradiction. This means that our initial assumption is false, and in reality:

$$\Big| \Pr[H_0 \to 1] - \Pr[H_1 \to 1] \Big| = \mathsf{negl}(n)$$

$\blacksquare$

3. **Claim 4.2** $\Pr[H_1 \to 1] = \frac{1}{2}$

Proof: Recall that in H_1, the challenge ciphertext is $c^* = (\mathsf{Enc}(k, 0^n), (r \oplus m_b))$. This ciphertext gives the adversary no information about b because m_b is masked by a uniformly random r.

The only place where r and b appear in the ciphertext is in $r \oplus m_b$. Furthermore, $r \oplus m_b$ is a random string that is independent of b because r is uniformly random. Therefore, the adversary's probability of guessing b correctly (i.e. the probability that $b' = b$) is exactly $\frac{1}{2}$. $\blacksquare$

4. Putting everything together, we have that:

$$\begin{aligned}
\Pr[H_0 \to 1] &\leq \Pr[H_1 \to 1] + \Big| \Pr[H_0 \to 1] - \Pr[H_1 \to 1] \Big| \\
&= \frac{1}{2} + \mathsf{negl}(n)
\end{aligned}$$

Since H_0 is the CPA security experiment for $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$, this means that $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ is CPA-secure.

A Flawed Approach

Here is a similar approach that doesn't quite work. It appeared on several student submissions, so it's useful to discus why it doesn't work.

Claim 4.3 *If* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is CPA-secure, then* $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ *is also CPA-secure.*

Proof:

1. Assume toward contradiction that $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ does not satisfy CPA security. Then there is an adversary $\mathcal{A}$ that wins the CPA security game for $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ with probability $\frac{1}{2}$+non-negligible(n). We will use $\mathcal{A}$ to construct an adversary $\mathcal{B}$ that breaks the CPA security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, which is a contradiction. Then we can conclude that $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ does satisfy CPA security.

2. Description of $\mathcal{B}$:

   (a) $\mathcal{B}$ runs $\mathcal{A}$ and responds to any queries made by $\mathcal{A}$. When $\mathcal{A}$ outputs a phase-I message m, $\mathcal{B}$ samples $r \leftarrow \{0,1\}^n$ and computes $\mathsf{Enc}(k, r)$ by making a phase-I query to its own oracle $\mathsf{Enc}(k, \cdot)$. Then $\mathcal{B}$ computes $c := \mathsf{Enc}(k, r), (r \oplus m)$ and sends c to $\mathcal{A}$.

   (b) In the challenge phase of $\mathcal{A}$'s CPA game, $\mathcal{A}$ outputs two messages (m_0, m_1). $\mathcal{B}$ samples a new $r \leftarrow \{0,1\}^n$ and outputs $(r, 0^n)$ as its challenge messages. $\mathcal{B}$ receives a ciphertext C^* from the challenger, which is either either $\mathsf{Enc}(k, r)$ or $\mathsf{Enc}(k, 0^n)$. $\mathcal{B}$ sends $(C^*, (r \oplus m_0))$ to $\mathcal{A}$.

   (c) In phase II of $\mathcal{B}$'s CPA game, $\mathcal{B}$ will simulate phase II of $\mathcal{A}$'s CPA game. This works the same way as phase I.

   (d) Finally, $\mathcal{A}$ will output a bit b', which $\mathcal{B}$ also outputs.

3. When $C^* = \mathsf{Enc}(k, r)$, then $\mathcal{A}$ receives $(\mathsf{Enc}(k, r), (r \oplus m_0))$, which is a valid encryption of m_0 under Enc_1. **In this case, the probability that $\mathcal{B}$ wins the CPA game is $\frac{1}{2} + \delta(n)$, where δ is some non-negligible function (*This analysis is incorrect*).**

4. When $C^* = \mathsf{Enc}(k, 0^n)$, then $\mathcal{A}$ receives $\mathsf{Enc}(k, 0^n), (r \oplus m_0)$, which is not a valid encryption of m_0. In fact, it gives no information about m_0 because m_0 is masked by a uniformly random string r. Therefore, in this case, the probability that $\mathcal{B}$ wins the CPA game is exactly $\frac{1}{2}$.

5. In summary, the probability that $\mathcal{B}$ wins the CPA game is:

$$\frac{1}{2} \cdot \left(\frac{1}{2} + \delta(n)\right) + \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{2} + \frac{\delta}{2}$$

   which is still non-negligibly greater than $\frac{1}{2}$.

■

To illustrate the problem with this analysis, let's construct an adversary $\mathcal{A}$ that is purposefully unhelpful to $\mathcal{B}$.

Description of $\mathcal{A}$: Let's say that $\mathcal{A}$ can decrypt any ciphertext generated by $\mathsf{Enc}_1(k, \cdot)$. When $\mathcal{A}$ receives its challenge ciphertext, let $\mathcal{A}$ decrypt the ciphertext to get m^*. If $m^* \neq m_1$, then $\mathcal{A}$ samples $b' \leftarrow \{0,1\}$ uniformly at random and outputs b'. If $m^* = m_1$, then $\mathcal{A}$ outputs $b' = 1$.

Next, the claims below show that $\mathcal{A}$ breaks the CPA security of $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$, but $\mathcal{B}$ does not break the CPA security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. Recall that the goal of the proof was to show that if $\mathcal{A}$ breaks the CPA security of $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$, then $\mathcal{B}$ breaks the CPA security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, so the proof is unsuccessful.

Claim 4.4 *$\mathcal{A}$ wins the CPA security game for $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ with probability $\frac{3}{4}$.*

Proof: If the CPA challenger encrypted m_0, then $\mathcal{A}$ outputs the correct answer ($b' = 0$) with probability $\frac{1}{2}$. If the CPA challenger encrypted m_1, then $\mathcal{A}$ outputs the correct answer ($b' = 1$) with probability 1. Therefore, $\mathcal{A}$ wins the CPA security game with probability

$$\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 1 = \frac{3}{4}$$

■

Claim 4.5 *$\mathcal{B}$ wins the CPA security game for $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with probability $\frac{1}{2}$.*

Proof: If the CPA challenger encrypted M_0, then $\mathcal{B}$ outputs the correct answer ($B' = 0$) with probability $\frac{1}{2}$. If the CPA challenger encrypted M_1, then $\mathcal{B}$ outputs the correct answer ($B' = 1$) with probability $\frac{1}{2}$. Therefore, $\mathcal{B}$ wins the CPA security game with probability $\frac{1}{2}$. ■

Another (valid) solution

This solution is a little simpler than the first solution given above. Rather than using $\mathcal{A}$ to distinguish between a valid and an invalid ciphertext, we will use $\mathcal{A}$ to distinguish between two valid ciphertexts. This solution is due to an anonymous student – thank-you to them!

Solution:

Claim 4.6 *If $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is CPA-secure, then $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ is also CPA-secure.*

Proof:

1. Assume toward contradiction that $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ does not satisfy CPA security. Then there is an adversary $\mathcal{A}$ that wins the CPA security game for $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ with probability $\frac{1}{2}$ plus non-negligible. We will use $\mathcal{A}$ to construct an adversary $\mathcal{B}$ that breaks the CPA security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, which is a contradiction. Then we can conclude that $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ does satisfy CPA security.

2. Description of $\mathcal{B}$:

   (a) $\mathcal{B}$ runs $\mathcal{A}$ and responds to any queries made by $\mathcal{A}$. When $\mathcal{A}$ outputs a phase-I message m, $\mathcal{B}$ samples $r \leftarrow \{0,1\}^n$ and computes $\mathsf{Enc}(k, r)$ by making a phase-I query to its own oracle $\mathsf{Enc}(k, \cdot)$. Then $\mathcal{B}$ computes $c := (\mathsf{Enc}(k, r), (r \oplus m))$ and sends c to $\mathcal{A}$.

   (b) In the challenge phase of $\mathcal{A}$'s CPA game, $\mathcal{A}$ outputs two messages (m_0, m_1). $\mathcal{B}$ samples a ciphertext $c_1^* \leftarrow \{0,1\}^n$. Then it computes $r_0 := c_1^* \oplus m_0$ and $r_1 := c_1^* \oplus m_1$. Then $\mathcal{B}$ outputs (r_0, r_1) as its challenge messages.

   The CPA challenger samples $B \leftarrow \{0,1\}$ and sends $\mathsf{Enc}(k, r_B)$ to $\mathcal{B}$. Then $\mathcal{B}$ sends

   $$(\mathsf{Enc}(k, r_B), c_1^*)$$

   to $\mathcal{A}$.

   (c) In phase II of $\mathcal{B}$'s CPA game, $\mathcal{B}$ will simulate phase II of $\mathcal{A}$'s CPA game. This works the same way as phase I.

   (d) Finally, $\mathcal{A}$ will output a bit b', which $\mathcal{B}$ also outputs.

3. Note that for either value of B, $(\mathsf{Enc}(k, r_B), c_1^*)$ is a valid encryption of m_B under $\mathsf{Enc}_1(k, \cdot)$. This is because $c_1^* = (r_B \oplus m_B)$. Furthermore, r_B is uniformly random and independent of (m_0, m_1, B). So $\mathcal{B}$ has correctly simulated the CPA security game for $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$.

4. If $\mathcal{A}$ wins the simulated CPA security game for $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$, by correctly guessing B, then $\mathcal{B}$ wins the CPA security game for $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. Therefore, $\mathcal{A}$'s success probability in the CPA security game for $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ equals $\mathcal{B}$'s success probability in the CPA security game for $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. If $\mathcal{A}$ breaks the CPA security of $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ (by winning the corresponding CPA game with probability $\frac{1}{2} + \text{non-negligible}(n)$), then $\mathcal{B}$ breaks the CPA security of $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$.

5. We know that $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is CPA-secure, so our initial assumption must be false, and in truth, $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ is also CPA-secure.

$\blacksquare$

Name:

5 Perfectly Secret Encryption (15 points)

In this problem we propose a definition of perfect secrecy for the encryption of *two* messages. We will prove that this definition *cannot* be satisfied by any encryption scheme.

Notation: We consider distributions over *pairs* of messages from the message space $\mathcal{M}$; we let M_1 and M_2 be random variables denoting the first and second message, respectively. (We stress that these random variables are not assumed to be independent.)

Encryption works as follows:

1. Generate a (single) key k, sample a pair of messages (m_1, m_2) according to the given distribution.
2. Compute ciphertexts $c_1 \leftarrow \mathsf{Enc}(k, m_1)$ and $c_2 \leftarrow \mathsf{Enc}(k, m_2)$; this induces a distribution over pairs of ciphertexts, and we let C_1 and C_2 be the corresponding random variables.

Proposed definition: Let us say that an encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is *perfectly secret for two messages* if for all distributions over $\mathcal{M} \times \mathcal{M}$, all $(m_1, m_2) \in \mathcal{M} \times \mathcal{M}$, and all ciphertexts $(c_1, c_2) \in \mathcal{C} \times \mathcal{C}$ for which $\Pr[C_1 = c_1 \wedge C_2 = c_2] > 0$,

$$\Pr[M_1 = m_1 \wedge M_2 = m_2 \mid C_1 = c_1 \wedge C_2 = c_2] = \Pr[M_1 = m_1 \wedge M_2 = m_2].$$

Question: Prove that no encryption scheme can satisfy the definition above. (You may assume that the encryption scheme satisfies perfect correctness).

Solution:

1. The key insight is that if $m_1 \neq m_2$, then the corresponding ciphertexts, c_1 and c_2, must not be equal. For any encryption scheme, the decryption function must be perfectly correct: every ciphertext should be decrypted to the correct message. If it were possible that $m_1 \neq m_2$ and $c := c_1 = c_2$, then the decryption of c would sometimes be incorrect.
2. Now, let's state that idea formally. Choose the distribution of the messages (M_1, M_2) such that $0 < \Pr[M_1 = M_2] < 1$. Then $0 < \Pr[C_1 = C_2]$ as well.
3. Choose a value $c \in \mathcal{C}$ such that $\Pr[C_1 = C_2 = c] > 0$. Then set $c_1 = c_2 = c$.
4. Choose m_1 and m_2 such that $m_1 \neq m_2$, and $\Pr[M_1 = m_1 \wedge M_2 = m_2] > 0$.
5. By the correctness of the encryption scheme,

   $$\Pr[C_1 = C_2 = c \mid M_1 = m_1 \wedge M_2 = m_2] = 0$$

   which implies, by Bayes' theorem, that

   $$\Pr[M_1 = m_1 \wedge M_2 = m_2 \mid C_1 = C_2 = c] = 0$$

6. However $\Pr[M_1 = m_1 \wedge M_2 = m_2] > 0$. Therefore:

$$\Pr[M_1 = m_1 \wedge M_2 = m_2 \mid C_1 = C_2 = c] \neq \Pr[M_1 = m_1 \wedge M_2 = m_2]$$

   This means that the scheme does not satisfy the definition given above for perfect secrecy for two messages.