


Q1 IND-CPA

(5 points)

When formalizing the notion of confidentiality, as provided by a proposed encryption scheme, we introduce the concept of indistinguishability under a chosen plaintext attack, or IND-CPA security. A scheme is considered *IND-CPA secure* if an attacker cannot gain any information about a message given its ciphertext. This definition can be defined as an experiment between a challenger and adversary, detailed in the diagram below:



Consider the one-time pad encryption scheme discussed in class. For parts (a) – (c), we will prove why one-time pad is not IND-CPA secure and, thus, why a key should not be reused for one-time pad encryption.

Q1.1 (1 point) What messages (M_0 and M_1) should the adversary provide the challenger?

Solution: The adversary can provide any two plaintexts A and B of same length to be encrypted.

Q1.2 (1 point) Now, for which message(s) should the adversary request an encryption from the challenger during the query phase?

Solution: The adversary can request an encryption for A , B , or both. Note that the adversary can request an arbitrary number of plaintexts to be encrypted, and can request the encryption of the same messages provided in the challenge phase.



(Question 1 continued...)

Q1.3 (1 point) The challenger now randomly selects $b \in \{0, 1\}$, encrypts M_b , and sends back $C = \text{Enc}(k, M_b) = M_b \oplus K$ to the adversary. How can the adversary find b with probability $> \frac{1}{2}$?

Solution: Since one-time pad is a deterministic encryption scheme, the ciphertext C we receive from the challenger will be identical to one of the ciphertexts we receive in the query phase. The adversary can simply compare C to $\text{Enc}(A)$ and $\text{Enc}(B)$ received in the query phase to determine which message was encrypted with probability 1.

Q1.4 (1 point) Putting it all together, explain how an adversary can always win the IND-CPA game with a probability 1 against a deterministic encryption algorithm. *Note: Given an identical plaintext, a deterministic encryption algorithm will produce identical ciphertext.*

Solution: An adversary can provide two plaintexts A and B to be encrypted. The adversary gets back X , which is an encryption of either A or B . Then, the adversary requests an encryption of A again and compares it with X . If the two are the same, X is the encryption of A , and vice versa.

Q1.5 (1 point) Assume that an adversary chooses an algorithm and runs the IND-CPA game a large number of times, winning with a probability of 0.6. Is the encryption scheme IND-CPA secure? Why or why not?

☐ Secure ☒ Insecure

Solution: By definition, a scheme is IND-CPA secure if the adversary wins with a probability $0.5 \pm \epsilon$, where ϵ is a negligibly small number.

In this case, the adversary has a non-negligible advantage in the IND-CPA game.


Q2 Block Ciphers I

(5 points)

Consider the Cipher Feedback (CFB) mode, whose encryption is given as follows:

$$C_i = \begin{cases} \text{IV}, & \text{if } i = 0 \\ E_K(C_{i-1}) \oplus P_i, & \text{otherwise} \end{cases}$$

Q2.1 (1 point) Draw the encryption diagram for CFB mode.



Q2.2 (1 point) What is the decryption formula for CFB mode?

$$P_i = E_K(C_{i-1}) \oplus C_i$$

Q2.3 (1 point) Select the true statements about CFB mode:

- ☐ Encryption can be parallelized
 ☒ The scheme is IND-CPA secure
- ☒ Decryption can be parallelized
 ☐ None of the above

Solution: Encryption is not parallelizable because encryption of block n of the plaintext is dependent on block $n - 1$ of the ciphertext.

Decryption is parallelizable because the decryption of block n of the ciphertext is dependent only on block $n - 1$ of the ciphertext.

The scheme is IND-CPA secure because an adversary cannot provide two messages of equal length such that they gain a non-negligible advantage in the IND-CPA game, as long as the IV is not reused. Note that if the IV is reused, the scheme would be deterministic.

(Question 2 continued...)

Q2.4 (1 point) What happens if two messages are encrypted with the same key and IV? What can the attacker learn about the two messages just by looking at their ciphertexts?

Solution: If the IV is reused in AES-CFB, the attacker can determine if two messages have identical prefix, up to but not including the first block containing the difference. This is because block n of the plaintext affects the input to the $n + 1$ st block cipher, so any difference in the plaintext will result in a completely different block cipher output for all future values.

When we use non-repeating IVs for CFB mode, even if we encrypt the same message multiple times, CFB-mode will generate distinct and random-looking ciphertexts each time.

Q2.5 (1 point) If an attacker recovers the IV used for a given encryption, but not the key, will they be able to decrypt a ciphertext encrypted with the recovered IV and a secret key? Explain why or why not.

☐ Yes ☒ No

Solution: The secrecy of the IV does not affect the security of the encryption scheme, as the IV is passed as part of the output of an encryption. The only condition is that the IV must not be reused in order for the given scheme to be secure.


Q3 Block Ciphers 3

(3 points)

Consider the following block cipher mode of operation.

M_i is the i -th block of plaintext. C_i is the i -th block of ciphertext. E_K is AES encryption with key K .

$$C_0 = M_0 = \text{IV} \quad C_i = E_K(M_{i-1} \oplus M_i)$$



Q3.1 (1 point) Which of the following is true about this scheme? Select all that apply.

- ☒ The encryption algorithm is parallelizable
- ☐ If one byte of a plaintext block M_i is changed, then the corresponding ciphertext block C_i will be different in exactly one byte.
- ☐ If one byte of a plaintext block M_i is changed, then the next ciphertext block C_{i+1} will be different in exactly one byte
- ☒ The encryption algorithm requires padding the plaintext
- ☐ None of the above

Solution: By looking at the equation or diagram, we can see that the ciphertext block C_i does not depend on any previous ciphertext block (it only depends on plaintext blocks M_{i-1} and M_i).

Since the plaintext block is passed through a block cipher, changing one byte of block cipher input will cause the block cipher output to be completely different

Changing one byte of M_i will change one byte of $M_i \oplus M_{i+1}$, the input to the block cipher. Again, changing one byte of block cipher input will cause the block cipher output to be completely different.

Since the plaintext block is XOR'd with the previous block of plaintext before being passed into a block cipher, the corresponding ciphertext blocks are not necessarily identical.

The plaintext is passed as an input to the block cipher, so it must be padded to a multiple of the block size.

(Question 3 continued...)

Q3.2 (1 point) TRUE OR FALSE: If the IV is always a block of all 0's for every encryption, this scheme is IND-CPA secure. Briefly justify your answer.

☐ TRUE ☒ FALSE

Solution: There is no randomness, so the scheme must be deterministic, and deterministic schemes cannot be IND-CPA secure.

Q3.3 (1 point) TRUE OR FALSE: If the IV is randomly generated for every encryption, this scheme is IND-CPA secure. Justify your answer.

☐ TRUE ☒ FALSE

Solution: Informally, note that the randomness in the IV is not passed to subsequent blocks. The second block uses the second plaintext block M_2 and the previous plaintext block M_1 as block cipher input, but never uses the IV . This is the case for all subsequent blocks as well. As a result, this scheme still leaks the existence of identical blocks.

Formally, here are some ways Eve could win the IND-CPA game:

- Sending $M = X \parallel X \parallel X$ (ciphertext C) and $M' = X \parallel Y \parallel Z$ (ciphertext C') results in the last two blocks of C being identical.
- Sending $M = 0 \parallel X$ and $M' = Y \parallel X$ results in distinguishable ciphertexts.
- Sending the same message twice results in everything but the first block of the ciphertext being identical.