Ciphertext indistinguishability

Ciphertext indistinguishability is an important security property of many encryption schemes. Intuitively, if a cryptosystem possesses the property of indistinguishability, then an adversary will be unable to distinguish pairs of ciphertexts based on the message they encrypt. The property of indistinguishability under chosen plaintext attack is considered a basic requirement for most provably secure public key cryptosystems, though some schemes also provide indistinguishability under chosen ciphertext attack and adaptive chosen ciphertext attack. Indistinguishability under chosen plaintext attack is equivalent to the property of semantic security, and many cryptographic proofs use these definitions interchangeably.

A cryptosystem is considered "secure in terms of indistinguishability" if no adversary A, given an encryption of a message randomly chosen from a two-element message space determined by the adversary, can identify the message choice with probability significantly better than that of random guessing $(1/2)$. If any adversary can succeed in distinguishing the chosen ciphertext with a probability significantly greater than $1/2$, then this adversary is considered to have an "advantage" in distinguishing the ciphertext, and the scheme is "not" considered secure in terms of indistinguishability. This definition encompasses the notion that in a secure scheme, the adversary should glean no information from seeing a ciphertext. Therefore, the adversary should be able to do no better than if it guessed randomly.

Formal definitions

Security in terms of indistinguishability has many definitions, depending on assumptions made about the capabilities of the attacker. It is normally presented as a game, where the cryptosystem is considered secure if no adversary can win the game with significantly greater probability than an adversary who must guess randomly. The most common definitions used in cryptography are **indistinguishability under chosen plaintext attack** (abbreviated IND-CPA), **indistinguishability under (non-adaptive) chosen ciphertext attack** (IND-CCA), and **indistinguishability under adaptive chosen ciphertext attack** (IND-CCA2). Security under either of the latter definition implies security under the previous ones: a scheme which is IND-CCA secure is also IND-CPA secure, and a scheme which is IND-CCA2 secure is both IND-CCA and IND-CPA secure. Thus, IND-CCA2 is the strongest of the these three definitions of security.

Indistinguishability under chosen-plaintext attack (IND-CPA)

For a probabilistic asymmetric key encryption algorithm, indistinguishability under chosen plaintext attack (IND-CPA) is defined by the following game between an adversary and a challenger. For schemes based on computational security, the adversary is modeled by a probabilistic polynomial time Turing machine, meaning that it must complete the game and output a "guess" within a polynomial number of time steps. In this definition E(PK, "M") represents the encryption of a message "M" under the key "PK":

The challenger generates a key pair "PK", "SK" based on some security parameter "k" (e.g., a key size in bits), and publishes "PK" to the adversary. The challenger retains "SK".

The adversary may perform any number of encryptions or other operations.

Eventually, the adversary submits two distinct chosen plaintexts M_0, M_1 to the challenger.

The challenger selects a bit "b" in {0, 1} uniformly at random, and sends the "challenge" ciphertext "C" = E(PK, M_b) back to the adversary.

The adversary is free to perform any number of additional computations or encryptions. Finally, it outputs a guess for the value of "b".

A cryptosystem is **indistinguishable under chosen plaintext** attack if every probabilistic polynomial time adversary has only a negligible "advantage" over random guessing. An adversary is said to have a negligible "advantage" if it wins the above game with probability (1/2) + epsilon(k), where epsilon(k) is a negligible function in the security parameter "k", that is for every (nonzero) polynomial function poly() there exists k_0 such that $|epsilon(k)|$ <|1/poly(k)| for all k >k_0.

Although the adversary knows M_0, M_1 and PK, the probabilistic nature of E means that the encryption of M_b will be only one of many valid ciphertexts, and therefore encrypting M_0, M_1 and comparing the resulting ciphertexts with the challenge ciphertext does not afford any advantage to the adversary.

While the above definition is specific to an asymmetric key cryptosystem, it can be adapted to the symmetric case by replacing the public key encryption function with an "encryption oracle", which retains the secret encryption key and encrypts arbitrary ciphertexts at the adversary's request.

Indistinguishability under chosen ciphertext attack/adaptive chosen ciphertext attack (IND-CCA, IND-CCA2)

Indistinguishability under non-adaptive and adaptive Chosen Ciphertext Attack (IND-CCA, IND-CCA2) uses a definition similar to that of IND-CPA. However, in addition to the public key (or encryption oracle, in the symmetric case), the adversary is given access to a "decryption oracle" which decrypts arbitrary ciphertexts at the adversary's request, returning the plaintext. In the non-adaptive definition, the adversary is allowed to query this oracle only up until it receives the challenge ciphertext. In the adaptive definition, the adversary may continue to query the decryption oracle even after it has received a challenge ciphertext, with the caveat that it may not pass the challenge ciphertext for decryption (otherwise, the definition would be trivial).

The challenger generates a key pair "PK", "SK" based on some security parameter "k" (e.g., a key size in bits), and publishes "PK" to the adversary. The challenger retains "SK".

The adversary may perform any number of encryptions, calls to the decryption oracle based on arbitrary ciphertexts, or other operations.

Eventually, the adversary submits two distinct chosen plaintexts M_0, M_1 to the challenger.

The challenger selects a bit "b" ϵ {0, 1} uniformly at random, and sends the "challenge" ciphertext "C" = E(PK, M_b) back to the adversary.

The adversary is free to perform any number of additional computations or encryptions.

In the "non-adaptive" case (IND-CCA), the adversary may "not" make further calls to the decryption oracle.

In the "adaptive" case (IND-CCA2), the adversary may make further calls to the decryption oracle, but may not submit the challenge ciphertext "C".

Finally, the adversary outputs a guess for the value of "b". A scheme is IND-CCA/IND-CCA2 secure if no adversary has a non-negligible advantage in winning the above game.

Non-malleability

Malleability is the property of some cryptographic algorithm . An encryption algorithm is malleable if it is possible for an adversary to transform a cipher-text into another ciphertext which decrypts to a related plaintext .That is given an encryption of a plaintext m, it is possible to generate another ciphertext which decrypts to f(m), for a known function f, without necessarily knowing or learning m.

Malleability is often an undesirable property in a general-purpose cryptosystem, since it allows an attacker to modify the contents of a message. For example, suppose that a bank uses a stream cipher to hide its financial information, and a user sends an encrypted message containing, say, "TRANSFER Rs.1000 TO ACCOUNT 199." If an attacker can modify the message on the wire, and can guess the format of the unencrypted message, the attacker could be able to change the amount of the transaction, or the recipient of the funds, e.g. "TRANSFER Rs1000 TO ACCOUNT 227."

Non-malleability is that given the ciphertext it is impossible to generate a different ciphertext so that the respective plane texts are related .

Motivation to follow non-malleability in cyptographir algorithm :

A well-established, albeit implicit, notion of non-malleability is existential unforgeability of signature schemes ;. Informally, a signature scheme is existentially unforgeable if, given access to ((m1; S(m1)); : : : ; (mk; S(mk)), where S(mi) denotes a signature on message mi, the adversary cannot construct a single valid (m; S(m)) pair for any new message m - even a nonsense message or a function of m1; : : : ; mk. Thus, existential unforgeability for signature schemes is the "moral equivalent" of non-malleability for cryptography.

Non-malleability is also important in private-key cryptography. Many common protocols, such as Kerberos or the Andrew Secure Handshake, use private key encryption as a sort of authentication mechanism: parties A and B share a key KAB. A sends to B the encryption of a nonce N under KAB, and the protocol requires B to respond with the encryption under KAB of f (N), where f is some simple function such as f (x) = x $\backslash Gamma$ The unproved and unstated assumption is that seeing KAB(N) doesn't help an imposter falsely claiming to be B to compute KAB(f (N)). As we shall see, this is precisely the guarantee provided by non-malleability.

Some Mathematical Deduction :

• A public-key scheme (E,D,G) is (t,q,ϵ)-secure in the NM-X sense if for all message distributions M, and all relations R:MxM 0,1, and for every

adversary A that runs in time t, and makes at most q queries to oracle O, there exists another adversary A' that runs in time poly(t), such that

Pr(pk , sk),m [R (m, Dsk (AO (pk , E pk (m)))] – Pr(pk , sk), m [R (m, Dsk (A '(pk)))] $\leq \epsilon$ (n)

where the oracle is:

$$O = \begin{cases} \text{-, if IND-CPA} \\ \text{Dsk , if IND-CCA} \end{cases}$$

and the adversary cannot query the decryption oracle at E pk (m)

• If a public-key scheme is (t,q,ϵ)-secure in NM-X sense, then it is (t,q,2ϵ)- secure in IND-X sense.

• Contradict that the scheme is (t,q,2ϵ)- secure in IND-X sense.

• Show that the scheme is also not (t,q,ϵ)- secure in NM-X sense.