

Lecture 2 - Perfect Secrecy and its Limitations

Boaz Barak

September 20, 2007

How do we define an “unbreakable encryption”?

Simplest case: shared key, encryption of single message.

Some notations: k - key of length $|k| = n$. (chosen at random in $\{0, 1\}^n$)

Plaintext: x of length $|x| = m$

Ciphertext: y

$y = E_k(x), x = D_k(y)$

$E : \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^\ell$

$D : \{0, 1\}^\ell \times \{0, 1\}^m \rightarrow \{0, 1\}^m$

Recurring themes in such definitions: probability, quantifiers

Motivating scenario Attacker knows that message is either **attack** or **retreat**.

Definition Two probability distributions X, Y over $\{0, 1\}^\ell$ are *identical*, denoted $X \equiv Y$, if for every $y \in \{0, 1\}^\ell$, $\Pr[X = y] = \Pr[Y = y]$.

The scheme (D, E) is *perfectly secure* if for every pair of messages x, x' , $E_{U_n}(x) \equiv E_{U_n}(x')$.

Exercise Does this mean that for every k , $E_k(x) = E_k(x')$?

Meaning If the message was **attack** then eavesdropper to see a ciphertext y sampled from $E_{U_n}(\text{attack})$. If the message was **retreat** then eavesdropper gets a ciphertext sampled from $E_{U_n}(\text{retreat})$. But this is the same distribution in both cases!

Game view of definition

- Adv chooses $x_1, x_2 \in \{0, 1\}^m$
- Sender selects $k \leftarrow_{\text{R}} \{0, 1\}^n$, $i \leftarrow_{\text{R}} \{1, 2\}$ and gives adversary $y = E_k(x_i)$.
- Adversary sends a number $j \in \{1, 2\}$.

Adversary *successful* if $i = j$.

Theorem 1. (E, D) is perfectly secure if and only if $\Pr[\text{Adv successful}] \leq 1/2$.

Proof.

If scheme is perfectly secure then adversary gets no information on i in this game and can win with probability at most $1/2$.

On the other hand, if scheme is not perfectly secure then there exist x_1, x_2 such that Y_{x_1} is not the same distribution as Y_{x_2} . This means that there is a string y_0 such that (without loss of generality) $\Pr[Y_{x_1} = y_0] > \Pr[Y_{x_2} = y_0]$.

Thus, if adversary does the following: pick x_1, x_2 and given ciphertext y , if $y = y_0$ then output 1, else output a random number in $\{1, 2\}$, then it will succeed with probability a little bit larger than $1/2$. \square

A perfectly secure encryption. The XOR operation takes two bits a and b and returns $a + b \pmod{2}$. Notation $a \oplus b$.

Some facts:

1. $a \oplus 0 = a$
2. $a \oplus a = 0$
3. Commutativity $a \oplus b = b \oplus a$
4. Associativity $a \oplus (b \oplus c) = (a \oplus b) \oplus c$

Generalize to vectors: $x, y \in \text{bits}^n$ then $x \oplus y = (x_1 \oplus y_1, \dots, z_b \oplus y_n)$.

The One-Time-Pad scheme: $n = |k| = |x|$, $E : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$.

$$E_k(x) = x \oplus k.$$

$$D_k(y) = y \oplus k$$

$$\text{Validity: } D_k(E_k(x)) = (x \oplus k) \oplus k = x \oplus (k \oplus k) = x \oplus 0^n = x$$

One-Time-Pad is perfectly secret. We'll prove that for every $x \in \{0, 1\}^n$, the distribution $Y_x = E_{U_n}(x)$ is distributed according to the uniform distribution. This means that that all these distributions are identical thus achieving perfect secrecy.

Proof. Let $y \in \{0, 1\}^n$, we need to show that $\Pr[Y_x = y] = 2^{-n}$, or in other words

$$\Pr_{k \leftarrow_R \{0, 1\}^n} [x \oplus k = y] = 2^{-n}$$

If we XOR x on the left side of both equations we get that $x \oplus x \oplus k = x \oplus y$ or in other words $k = x \oplus y$. That is, there a unique single value of k ($x \oplus y$) that will make the equation true. Since there is a total number of 2^n possible strings k , the probability the equation is true is 2^{-n} . \square

What about three messages? Suppose that the possible messages are **attack**, **retreat**, and **stay**. Is perfect security still good in this case?

Theorem 2. *If (E, D) is perfectly secure, then in the game where adversary picks three messages x_1, x_2, x_3 and then gets $E_k(x_i)$ for random k and $i \leftarrow_R \{1, 2, 3\}$ then adversary can guess i with probability at most $1/3$.*

Proof. The adversary gets the same output no matter what i was. So it gets no information on i . \square

Is this the end of cryptography? We have a simply, efficient provably unbreakable encryption scheme. What more do we want?

- Use the same key for many messages
- At a minimum use n bit key for $2n$ -bit messages.

In the exercise you will show that natural extension of one-time-pad to use shorter keys (the “two-time-pad”) is not secure.

Limitations of perfect security

Theorem 3. *There is no perfectly secret encryption scheme (E, D) with n -bit messages and $n - 1$ -bit keys.*

Proof. Suppose that (E, D) is such an encryption scheme. Denote by Y_0 the distribution $E_{U_{n-1}}(0^n)$ and by S_0 its support. Since there are only 2^{n-1} possible keys, $|S_0| \leq 2^{n-1}$.

Now for every key k the function $E_k(\cdot)$ is one to one and hence its image is of size at least 2^n . This means that for every k there exists x such that $E_k(x) \notin S_0$. Fix such a k and x , then the distribution $E_{U_{n-1}}(x)$ does not have the same support as Y_0 and hence it is not identical to it. \square

Objections to impossibility result Whenever faced with an impossibility result that says we can not do something we want, it is a good idea to examine the underlying assumptions behind this result, and see if we can relax these assumption to still get what we want (or at least something close to that).

Statistical security Suppose that we allowed the adversary to have a tiny advantage in its posteriori guessing probability compared to the a-priori probability.

For example, we can say that a scheme is ϵ *statistically indistinguishable* if the probability that adversary guesses which of the two messages was encrypted is at most $\frac{1}{2} + \epsilon$.

Whenever making a relaxation to a definition, two questions arise: **(1)** Is the relaxed definition still strong enough for applications? **(2)** Does the relaxation buy something we could not get with the original definition?

In this case, the answer to **(1)** is YES: If ϵ is very small (say 10^{-6} or maybe even 10^{-100}) then ϵ -secure schemes will be just as good as perfectly secure schemes for all practical purposes. Thus, if we could bypass the impossibility result for perfect secrecy using this notion this would be great.

Unfortunately, we will see that the answer to **(2)** is NO - we can't have key shorter than the message even if we relax to statistical security.

Equivalence It turns out this is again essentially equivalent to a relaxation of perfect secrecy which we call ϵ -*statistical secrecy*. This means that for every x, x' the distributions $Y_x = E_{U_n}(x)$ and $Y_{x'} = E_{U_n}(x')$, even if not identical, are still within at most ϵ *statistical distance*.

Definition 1. Let X and Y be two distributions over $\{0, 1\}^n$. The *statistical distance* of X and Y , denoted by $\Delta(X, Y)$ is defined to be

$$\max_{T \subseteq \{0, 1\}^n} \left| \Pr[X \in T] - \Pr[Y \in T] \right|$$

if $\Delta(X, Y) \leq \epsilon$ we say that $X \equiv_\epsilon Y$

Lemma 1.

$$\Delta(X, Y) = \frac{1}{2} \sum_{w \in \text{Supp}(X) \cup \text{Supp}(Y)} |\Pr[X = w] - \Pr[Y = w]|$$

Proof. left as exercise. □

Minimal key size for statistically secure schemes. Unfortunately, statistical security doesn't enable us to get much shorter keys. In fact, if we use just one bit less the adversary can get an advantage of $1/4$ (which we consider huge — remember that we were hoping for $10^{-100!}$). We'll prove the following theorem:

Theorem 4. *Let (E, D) be a valid encryption with $E : \{0, 1\}^n \times \{0, 1\}^{n+1} \rightarrow \{0, 1\}^*$. Then there exist plaintexts x_1, x_2 with $\Delta(x_1, x_2) > 0.1$.*

Proof. In the proof we'll use the following seemingly trivial observation: for a random variable Y , if $\mathbb{E}[Y] \leq \mu$ then $\Pr[Y \leq \mu] > 0$.

Let $x_1 = 0^{n+1}$ and let $S = \text{Supp}(E_{U_n}(x_1))$. Note that $|S| \leq 2^n$.

Consider the following experiment: we choose a random message $x \leftarrow_{\text{R}} \{0, 1\}^{n+1}$ and define the following 2^n random variables: for every k , $T_k(x) = 1$ if $E_k(x) \in S$ and 0 otherwise.

For every k , $E_k(\cdot)$ is one to one and hence $\Pr[T_k = 1] \leq 1/2$. This means that $\mathbb{E}[T_k] \leq 1/2$.

Define $T = \sum_{k \in \{0, 1\}^n} T_k$. Then

$$\mathbb{E}[T] = \mathbb{E}\left[\sum_k T_k\right] = \sum_k \mathbb{E}[T_k] \leq 2^n/2$$

This means that the probability that $\Pr[T \leq 2^n/2] > 0$ or in other words, there exists x such that $\sum_k T_k(x) \leq 2^n/2$. This means that for such x , at most half of the keys k satisfy $E_k(x) \in S$, or equivalently $\Pr[E_{U_n}(x) \in S] \leq 1/2$. Since $\Pr[E_{U_n}(0^{n+1}) \in S] = 1$ we get that

$$\Delta(E_{U_n}(0^{n+1}), E_{U_n}(x)) \geq 1/2$$

This technique — proving the existence of an object with a particular property by proving that the probability the property is satisfied is positive — is called the *probabilistic method*. There's a beautiful book about it with this name by Alon and Spencer. □