Notes for Lecture 2

In which we encounter for the first time message indistinguishability and semantic security

In the last lecture we saw that

- all classical encryption schemes which allow the encryption of arbitrarily long messages have fatal flaws;
- it is possible to encrypt with perfect security using one-time pad, but the scheme can be used only once, and the key has to be as long as the message;
- if one wants perfect security, one needs a key as long as the total length of all messages that are going to be sent.

Our goal for the next few lectures will be to study schemes that allow the sending of messages that are essentially arbitrarily long, using a fixed key, and having a security that is essentially as good as the perfect security of one-time pad.

Today we introduce a notion of security (*semantic security*) that is extremely strong. When it is met *there is no point for an adversary to eavesdrop the channel*, regardless of what messages are being sent, of what she already knows about the message, and what goal she is trying to accomplish.

First, let us fix the model in which we are going to work. For the time being, we are going to be very modest, and we shall only try to construct an encryption scheme that, like one-time pad, is designed for only one use. We just want the key to be reasonably short and the message to be of reasonably large length.

We shall also restrict ourselves to *passive* adversaries, meaning that Eve is able to see the communication between Alice and Bob, but she cannot inject her own messages in the channel, and she cannot prevent messages from being delivered.

The definition of *correctness* for an *encryption scheme* is straightforward.

Definition 1 (Symmetric-Key Encryption Scheme – Finite case) *A symmetric-key encryption scheme with key-length k, plain-text length m and ciphertext-length c is is a pair of probabilistic algorithms (Enc, Dec), such that $Enc : \{0,1\}^k \times \{0,1\}^m \to \{0,1\}^c$, $Dec : \{0,1\}^k \times \{0,1\}^c \to \{0,1\}^m$, and for every key $K \in \{0,1\}^k$ and every message M,*

$$\mathbb{P}[Dec(K, Enc(K, M)) = M] = 1 \tag{1}$$

where the probability is taken over the randomness of the algorithms

Definition 2 (Symmetric-Key Encryption Scheme – Variable Key Length Case) *A symmetric-key encryption scheme with variable key-length is is a pair of polynomial-time probabilistic algorithms* (Enc, Dec) *and a function* $m(k) > k$*, such that for every* security parameter k*, for every key* $K \in \{0,1\}^k$ *and every message* $M \in \{0,1\}^{m(k)}$*,*

$$\mathbb{P}[Dec(k, K, Enc(k, K, M)) = M] = 1 \tag{2}$$

(Although it is redundant to give the algorithm the parameter k, we do so because this will emphasize the similarity with the public-key setting that we shall study later.)

It will be more tricky to satisfactorily formalize the notion of *security*.

One super-strong notion of security, which is true for the one-time pad is the following:

Definition 3 (Perfect Security) *A symmetric-key encryption scheme* (Enc, Dec) *is* perfectly secure *if, for every two messages* M, M'*, the distributions* $Enc(K, M)$ *and* $Enc(K, M')$ *are identical, where we consider the distribution over the randomness of the algorithm* $Enc()$ *and over the choice of* $K \sim \{0,1\}^k$*.*

The informal discussion from the previous lecture gives a hint to how to solve the following

Exercise 1 *Prove that if* (Enc, Dec) *is perfectly secure, then* $k \geq m$*.*

Before we move on, let us observe two limitations that will be present in any possible definitions of security involving a key of length k which is much smaller than the message length. Eve can always employ one of the following two trivial attacks:

1. In time 2^k, Eve can enumerate all keys and produce a list of 2^k plaintexts, one of which is correct. Further considerations can help her prune the list, and in certain attack models (which we shall consider later), she can figure out with near certainty which one is right, recover the key and totally break the system.
2. Eve can make a random guess of what the key is, and be correct with probability 2^{-k}.

Already if $k = 128$, however, neither line of attack is worrisome for Alice and Bob. Even if Eve has access to the fastest of super-computers, Alice and Bob will be long dead of old age before Eve is done with the enumeration of all 2^{128} keys; and Alice and Bob are going to both going to be stricken by lighting, and then both hit by meteors, with much higher probability than 2^{-128}.

The point, however, is that *any* definition of security will have to involve a bound on Eve's running time, and allow for a low probability of break. If the bound on Eve's

running time is enormous, and the bound on the probability of a break is minuscule, then the definition is as satisfactory as if the former was infinite and the latter was zero.

All the definitions that we shall consider involve a bound on the complexity of Eve, which means that we need to fix a model of computation to measure this complexity. We shall use (non-uniform) *circuit complexity* to measure the complexity of Eve, that is measure the number of gates in a boolean circuit implementing Eve's functionality.

If you are not familiar with circuit complexity, the following other convention is essentially equivalent: we measure the running time of Eve (for example on a RAM, a model of computation that captures the way standard computers work) and we *add the length of the program* that Eve is running. The reason is that, without this convention, we would never be able to talk about the complexity of computing finite functions. Every function of an 128-bit input, for example, is very efficiently computable by a program which is nothing but a series of 2^{128} if-then-elses.

Finally we come to our first definition of security:

Definition 4 (Message Indistinguishability – concrete version) *We say that an encryption scheme (Enc, Dec) is (t, ϵ) message indistinguishable if for every two messages M, M', and for every boolean function T of complexity $\leq t$, we have*

$$|\mathbb{P}[T(Enc(K, M) = 1] - \mathbb{P}[T(Enc(K, M')) = 1]| \leq \epsilon \tag{3}$$

where the probability is taken over the randomness of $Enc()$ and the choice of $K \sim \{0,1\}^k$.

(Typical parameters that are considered in practice are $t = 2^{80}$ and $\epsilon = 2^{-60}$.)

When we have a family of ciphers that allow varying key lengths, the following asymptotic definition is standard.

Definition 5 (Negligible functions) *A function $\nu : \mathbb{N} \to \mathbb{R}^+$ is* negligible *if for every polynomial p and for every sufficiently large n*

$$\nu(n) \leq \frac{1}{p(n)}$$

Definition 6 (Message Indistinguishability – asymptotic definition) *We say that a variable key length encryption scheme (Enc, Dec) is message indistinguishable if for every polynomial p there is a negligible function ν such that for all sufficiently large k the scheme is $(p(k), \nu(k))$-message indistinguishable when the security parameter is k.*

The motivation for the asymptotic definition, is that we take polynomial time to be an upper bound to the amount of steps that any efficient computation can take, and to the "number of events" that can take place. This is why we bound Eve's running time by a polynomial. The motivation for the definition of negligible functions is that if an event happens with negligible probability, then the expected number of experiments that it takes for the event to happen is superpolynomial, so it will "never" happen. Of course, in practice, we would want the security parameters of a variable-key scheme to be exponential, rather than merely super-polynomial.

Why do we use message indistinguishability as a formalization of security?

A first observation is that if we take $\epsilon = 0$ and put no limit on t, then message-indistinguishability becomes perfect security, so at least we are dealing with a notion whose "limit" is perfect security.

A more convincing explanation is that message indistinguishability is equivalent to *semantic security*, a notion that we describe below and that, intuitively, says that Eve might as well not look at the channel.

What does it mean that "Eve might as well not look at the channel"? Let us summarize Eve's information and goals. Alice has a message M that is sent to Bob over the channel. The message comes from some distribution X (for example, it is written in English, in a certain style, it is about a certain subject, and so on), and let's assume that Eve knows X. Eve might also know more about the specific message being sent, because of a variety of reasons; call $I(M)$ the information that Eve has about the message. Finally, Eve has a goal in eavesdropping the conversation, which is to learn some information $f(M)$ about the message. Perhaps she wants to reconstruct the message in its entirety, but it could also be that she is only interested in a single bit. (Does M contain the string "I hate Eve"? Is M a confidential report stating that company Y is going to miss its earning estimate? And so on.)

Why is Eve even bothering tapping the channel? Because via some cryptoanalytic algorithm A, which runs in a reasonable amount of time, she thinks she has a good chance to accomplish. But the probability of accomplishing her goal would have been essentially the same *without* tapping the channel, then there is no point.

Definition 7 (Semantic Security – Concrete definition) *An encryption scheme (Enc, Dec) is (t, o, ϵ) semantically secure if for every distribution X over messages, every functions $I : \{0,1\}^m \to \{0,1\}^*$ and $f : \{0,1\}^m \to \{0,1\}^*$ (of arbitrary complexity) and every function A of complexity $t_A \leq t$, there is a function A' of complexity $\leq t_A + o$ such that*

$$|\mathbb{P}[A(Enc(K, M), I(m)) = f(M)] - \mathbb{P}[A'(I(m)) = f(M)]| \leq \epsilon$$

Think, as before, of $t = 2^{80}$ and $\epsilon = 2^{-60}$, and suppose o is quite small (so that a

computation of complexity o can be performed in a few seconds or less), and notice how the above definition captures the previous informal discussion.

Now let's see that semantic security is *equivalent* to message indistinguishability.

Lemma 8 (Semantic Security Implies Message Indistinguishability) *If (Enc, Dec) is (t, o, ϵ) semantically secure, then it is $(t, 2\epsilon)$ message indistinguishable.*

Note that semantic security implies message indistinguishability *regardless of the overhead parameter o.*

PROOF: We prove that if (Enc, Dec) is *not* $(t, 2\epsilon)$ message indistinguishable then it is *not* (t, o, ϵ) semantically secure regardless of how large is o.

If (Enc, Dec) is not $(t, 2\epsilon)$ message indistinguishable, then there are two messages M_0, M_1 and an algorithm T of complexity $\leq t$ such that

$$\mathbb{P}[T(Enc(K, M_1)) = 1] - \mathbb{P}[T(Enc(K, M_0)) = 1] > 2\epsilon \tag{4}$$

Pick a bit b uniformly at random in $\{0, 1\}$; then we have

$$\mathbb{P}[T(Enc(K, M_b)) = b] > \frac{1}{2} + \epsilon \tag{5}$$

And now take A to be T, X to be the distribution M_b for a random b, and define $f(M_b) = b$, and $I(M)$ to be empty. Then

$$\mathbb{P}[A(I(M), Enc(K, M)) = f(M)] > \frac{1}{2} + \epsilon \tag{6}$$

On the other hand, for every A', regardless of complexity

$$\mathbb{P}[A'(I(M), Enc(K, M)) = f(M)] = \frac{1}{2} \tag{7}$$

and so we contradict semantic security. □

Lemma 9 (Message Indistinguishability Implies Semantic Security) *If (Enc, Dec) is (t, ϵ) message indistinguishable and Enc has complexity $\leq p$, then (Enc, Dec) is $(t - \ell_f, p, \epsilon)$ semantically secure, where ℓ_f is the maximum length of $f(M)$ over $M \in \{0, 1\}^m$.*

PROOF: Fix a distribution X, an information function I, a goal function f, and a cryptoanalytic algorithm A of complexity $\leq t - \ell_f$.

Take $A'(I(M)) = A(I(M), Enc(K, \mathbf{0}))$, so that the complexity of A' is equal to the complexity of A plus the complexity of Enc.

For every message M, we have

$$\mathbb{P}[A(I(M), Enc(K, M)) = f(M)] \leq \mathbb{P}[A(I(M), Enc(K, \mathbf{0})) = f(M)] + \epsilon \quad (8)$$

Otherwise defining $T(C) = 1 \Leftrightarrow A(I(M), C) = f(M)$ would contradict the indistinguishabiltiy.

Averaging over M in X

$$\mathbb{P}_{M\sim X, K\in\{0,1\}^n}[A(I(M), Enc(K, M)) = f(M)] \leq \mathbb{P}_{M\sim X, K\in\{0,1\}^n}[A(I(M), Enc(K, \mathbf{0})) = f(M)] + \epsilon \quad (9)$$

and so

$$\mathbb{P}_{M\sim X, K\in\{0,1\}^n}[A(I(M), Enc(K, M)) = f(M)] \leq \mathbb{P}_{M\sim X, K\in\{0,1\}^n}[A'(I(M)) = f(M)] + \epsilon \quad (10)$$

□

It is also possible to define an asymptotic version of semantic security, and to show that it is equivalent to the asymptotic version of message indistinguishability.

Definition 10 (Semantic Security – Asymptotic Definition) *An encryption scheme (Enc, Dec) is semantically secure if for every polynomial p there exists a polynomial q and a negligible function ν such that (Enc, Dec) is $(p(k), q(k), \nu(k))$ semantically secure for all sufficiently large k.*

Exercise 2 *Prove that a variable key-length encryption scheme (Enc, Dec) is asymptotically semantically secure if and only if it is asymptotically message indistinguishable,*