University of Maryland
CMSC456 — Introduction to Cryptography
Professor Jonathan Katz

Lecture 3

1 The One-Time Pad

1.1 Proof of Security for the One-Time Pad

Recall the definition of perfect security (or secrecy) we had last time:

Definition 1 *An encryption scheme over message space $\mathcal{M}$ is* perfectly secure *if, for all distributions over $\mathcal{M}$, for all $m \in \mathcal{M}$, and for all ciphertexts c we have* $\Pr[m|c] = \Pr[m]$. *In other words, the* a postiori *probability that a message m was sent, given that we observe ciphertext c, is exactly equal to the* a priori *probability that message m is sent.*

We now give a full proof that the one-time pad encryption scheme is secure (last time we only gave a proof for the uniform distribution over $\mathcal{M}$).

Theorem 1 *The one-time pad is a perfectly-secure encryption scheme.*

Proof Assume $\mathcal{M} = \{0,1\}^n$. For any $m \in M$ and any c we have:

$$\begin{aligned} \Pr[m|c] &= \frac{\Pr[m \wedge c]}{\Pr[c]} \\ &= \frac{\Pr[c|m] \cdot \Pr[m]}{\Pr[c]}, \end{aligned} \tag{1}$$

using two applications of the definition of conditional probability. Conditioning over all messages gives $\Pr[c] = \sum_{m \in \mathcal{M}} \Pr[c|m] \cdot \Pr[m]$. But, for any m, c we have:

$$\begin{aligned} \Pr[c|m] &= \Pr[k = (c \oplus m)] \\ &= 2^{-n} \end{aligned}$$

so that $\Pr[c] = 2^{-n} \cdot \sum_{m \in \mathcal{M}} \Pr[m] = 2^{-n}$. Plugging into (1) shows that $\Pr[m|c] = \Pr[m]$ and we are done. ■

1.2 Optimality of the One-Time Pad

The one-time pad isn't a very good encryption scheme. For one thing, it cannot be used to send more than one message. Furthermore, you need to share n bits to send an n-bit message; but if you can meet in secret and agree on n bits, why not just meet in secret and hand over your message! A natural question is whether we can do better.

In fact, we cannot. The next theorem shows (roughly) that to perfectly encrypt n bits, you need to share at least n bits.

Theorem 2 *If $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ is a perfectly secure encryption scheme over message space $\mathcal{M}$, then we must have $|\mathcal{K}| \geq |\mathcal{M}|$ (or, roughly speaking, if $\mathcal{M} = \{0,1\}^n$ then we must have $|\mathcal{K}| \geq 2^n$ and the length of any particular key is n bits).*

Proof Say we observe ciphertext c. We can play the part of the receiver and decrypt c using every possible key $k \in \mathcal{K}$. This gives us at most $|\mathcal{K}|$ different messages which could possibly have resulted in ciphertext c (note: this argument holds even for a randomized encryption scheme, as long as we assume correctness of the decryption algorithm). But if $|\mathcal{K}| < |\mathcal{M}|$ then there is at least one message $m \in \mathcal{M}$ for which $\Pr[m|c] = 0$. Thus, the scheme will be insecure if the *a priori* probability of m is non-zero (which we can assure by choosing the distribution over $\mathcal{M}$ appropriately). ■

1.3 Stronger Attack Models

We mentioned earlier that the one-time pad is insecure if used twice (well, *obviously...*). We can rephrase this as follows. Imagine an arbitrary encryption scheme that is used to encrypt two messages from Alice to Bob (call these two messages m_1 and m_2). Certainly, it might happen that an eavesdropper knows what m_1 is (or at least, has some information about m_1): for example, m_1 might be an ACK message, or might be in English, or might represent a yes/no answer. A property we might desire from our encryption scheme is that *even if* the adversary knows m_1 and sees c_1, the encryption of m_2 should remain secure (i.e., observing c_2 should give no information about m_2). Note that, although reasonable, this is not the case for the one-time pad. If the adversary knows m_1 and then sees c_1, the adversary can immediately compute the key as $k = m_1 \oplus c_1$. Now any future ciphertexts that are observed by the adversary can be decrypted immediately!

Informally, then, we can define security against *known plaintext attacks* as follows (we will give a formal definition in a few weeks):

Definition 2 *A scheme is* secure against known plaintext attacks *if it is secure even when the adversary is given a sequence of pairs $(m_1, c_1 = \mathcal{E}_k(m_1)), \ldots, (m_\ell, c_\ell = \mathcal{E}_k(m_\ell))$, where $m_1, \ldots, m_\ell$ are randomly chosen. (Note that the same key is used throughout, and this same key is used for the ciphertext observed by the adversary that it is trying to decrypt.)*

The basic level of security achieved by the one-time pad is often refered to as *security against ciphertext only attacks*. I.e., the adversary gets no plaintext/ciphertext pairs before being asked to "decrypt" a particular ciphertext.

We can imagine an even more insidious type of attack than the above: how about a *chosen plaintext attack* where the adversary gets to choose which plaintexts are encrypted by Alice. What might this correspond to in real life? Well, the adversary might control an application-level protocol that is feeding data to Alice to be encrypted. Or, the adversary might be able to impersonate Bob and thereby force (or otherwise cause) Alice to encrypt certain things. Again, it would be nice if an encryption scheme could be secure for *future* messages even under this sort of attack. Informally (we give a more formal definition later on in the course):

Definition 3 *A scheme is* secure against chosen plaintext attacks *if it is secure even when the adversary chooses messages $m_1, \ldots, m_\ell$ and then gets to see ciphertexts $c_1 = \mathcal{E}_k(c_1), \ldots, c_\ell$. (Again, the same key is used throughout, and this same key will be used for the "challenge" ciphertext observed by the adversary later on.)*

Imagine for a moment how you would construct a scheme secure against a single known plaintext attack (i.e., where the adversary gets to see a single pair $(m_1, c_1 = \mathcal{E}_k(m_1))$. A little thought shows that perfect security in this setting is difficult to achieve; in fact, we can "prove" a theorem of the following form (of course, we can't really prove such a theorem until we give a rigorous version of definition 2; furthermore, under some reasonable definitions security against a single known plaintext attack *is* possible):

Theorem 3 *No (stateless) encryption scheme can be perfectly secure against known plaintext attacks. This is true even if the encryption scheme is randomized.*

You are asked to prove a version of this theorem on the homework.

2 Where do we go From Here?

So far we have seen two negative results, informally summarized here:

- To perfectly encrypt n-bit messages, we need to share n-bit keys.
- We can never achieve perfect security against known plaintext attacks (or, for that matter, chosen plaintext attacks) using a stateless encryption scheme. A consequence of this is that we *cannot achieve perfect secrecy when encrypting multiple messages* unless we use a stateful encryption scheme.

I will just mention that, in general, we want to avoid stateful encryption since this causes problems in case the sender and receiver get "out of sync". Also, it is a general rule that the less state the better (avoids tying up memory).

So, do we just give up? Is this the end of cryptography as we know it?! Well, if the definition is too hard to achieve, let's just relax the definition...

Before we relax the definition, let's look at some alternate ways of presenting Defintion 1. The following is nice because of its simplicity:

Definition 4 *An encryption scheme over message space $\mathcal{M}$ is perfectly secure if, for the uniform distribution over any set $\{m_1, m_2\} \subset \mathcal{M}$ of two messages and for all ciphertexts c we have:* $\Pr[m_1|c] = \Pr[m_2|c]$. *I.e., if m_1 and m_2 are the only possible messages, and these are equiprobable a priori, then they remain equiprobable even after observing c.*

Since we use the term *perfect security*, we had better hope that Definitions 4 and 1 are equivalent. In fact, the following proof shows that they are:

Proof One direction is easy. Clearly, if a scheme satisfies Definition 1 then choosing the distribution over $\mathcal{M}$ as the uniform distribution over some two-message subset immediately gives Definition 4.

Now, assume we have a scheme satisfying Definition 4. Fix m_1, m_2, c. Since $\Pr[m_1|c] + \Pr[m_2|c] = 1$ we have $\Pr[m_1|c] = \Pr[m_2|c] = 1/2$ (where these probabilities are in the experiment appropriate to Definition 4). Now,

$$\begin{aligned}
\Pr[m_1|c] &= \frac{\Pr[c|m_1] \cdot \Pr[m_1]}{\Pr[c]} \\
&= \frac{1/2 \cdot \Pr[c|m_1]}{1/2 \cdot (\Pr[c|m_1] + \Pr[c|m_2])} \\
&= \frac{\Pr[c|m_1]}{\Pr[c|m_1] + \Pr[c|m_2]}.
\end{aligned}$$

Since this equation holds true for arbitrary m_1, m_2 we must have $\Pr[c|m_1] = \Pr[c|m_2]$ for *any* m_1, m_2.

Now, take an arbitrary distribution over $\mathcal{M}$ and arbitrary m, c. We have:

$$\begin{aligned}
\Pr[m|c] &= \frac{\Pr[c|m] \cdot \Pr[m]}{\Pr[c]} \\
&= \frac{\Pr[c|m] \cdot \Pr[m]}{\sum_{m' \in \mathcal{M}} \Pr[c|m'] \Pr[m']} \\
&= \Pr[m],
\end{aligned}$$

where we use the fact that $\Pr[c|m'] = \Pr[c|m]$ for all m', and the fact that $\sum_{m' \in \mathcal{M}} \Pr[m'] = 1$. But this is exactly as required by Definition 1. ■

We introduce one more equivalent definition before our main modification:

Definition 5 *An encryption scheme over message space $\mathcal{M}$ is perfectly secure if, for any two messages $m_1, m_2 \in \mathcal{M}$ and for* any *algorithm A we have:*

$$\Pr[A(C) = m_1 | C = \mathcal{E}_k(m_1)] = \Pr[A(C) = m_1 | C = \mathcal{E}_k(m_2)].$$

A word about this definition. What we have is an adversary (algorithm) A who is trying to "guess" which of two possible messages is being sent. The notation "$A(C)$" refers to the output of the algorithm when it is run on input C. Thus, we require that the probability that A guesses m_1 when it is given an encryption of m_1 should be exactly the same as the probability that A guesses m_1 when it is given an encryption of m_2. In other words, A is just randomly guessing and is not doing any better at guessing regardless of whether it is given an encryption of one message or the other.

It is not too hard (but it is a little tedious) to show that this definition is equivalent to the one before, and hence to Definition 1. Thus, it is as difficult to obtain as the original definition. But we now modify it in two crucial ways. First, we do not require security against *any* algorithm, but instead against *efficient* algorithms (we say more about this next lecture). Second, we do not require that the algorithm (adversary) have exactly equal probabilities of guessing the message in each case; we only require that they be close (we say more about how close next lecture). Thus:

Definition 6 *An encryption scheme over message space $\mathcal{M}$ is* secure *(however, no longer perfectly secure) if, for any two messages $m_1, m_2 \in \mathcal{M}$ and any* efficient *algorithm A we have:*

$$|\Pr[A(C) = m_1 | C = \mathcal{E}_k(m_1)] - \Pr[A(C) = m_1 | C = \mathcal{E}_k(m_2)]| < \epsilon.$$