Crypto-Scope 12/28/2009

Lecture Notes 2: Shannon's theory on Perfect Secrecy

Professor: Kannan Srinathan *Scribe:Rohit and Sandeep*

1 Concept

In the previous lecture, we saw classical encryption schemes and showed how these can be broken and ways to improve such ciphers. In this lecture, we will look into encryption schemes that are provably secure against an adversary which is computationally bounded. We begin by briefly recalling some of the syntax that was introduced in previous lecture. An encryption scheme is defined by three algorithms *Gen*, *Enc* and *Dec*, as well as a specification of a message space $\mathcal{M}$ with $|\mathcal{M}| > 1$. The key generation algorithm *Gen* is a probabilistic algorithm that outputs a key k chosen according to some distribution. We denote by $\mathcal{K}$ the key space, i.e., the set of all possible keys that can be output by *Gen*, and require $\mathcal{K}$ to be finite. The encryption algorithm *Enc* takes as input a key $k \in \mathcal{K}$ and a message $m \in \mathcal{M}$, and outputs a cipher text c; we denote this by $Enc_k(m)$. The encryption algorithm may be probabilistic, so that $Enc_k(m)$ might output a different cipher text when run multiple times. We let C denote the set of all possible cipher texts that can be output by $Enc_k(m)$, for all possible choices of $k \in \mathcal{K}$ and $m \in \mathcal{M}$.The decryption algorithm *Dec* takes as input a key $k \in \mathcal{K}$ and a cipher text $c \in C$ and outputs a message $m \in \mathcal{M}$. We write $Dec_k(c)$ to denote the process of decrypting cipher text c using key k.

1.1 Definition of Perfect Secrecy

Definition 1 *An encryption scheme (Gen,Enc,Dec) over a message space $\mathcal{M}$ is* perfectly secret *if for every probability distribution over $\mathcal{M}$, every message $M \in \mathcal{M}$, and every cipher text $c \in C$ for which $Pr[C = c] > 0$:*

$$Pr[M = m \mid C = c] = Pr[M = m].$$

Lemma 1 *An encryption scheme (Gen,Enc,Dec) over a message space $\mathcal{M}$ is* perfectly secret *if and only if for every probability distribution over $\mathcal{M}$, every message $m \in \mathcal{M}$, and every cipher text $c \in C$:*

$$Pr[C = c \mid M = m] = Pr[C = c].$$

Proof: Fix a distribution over $\mathcal{M}$ and arbitrary $m \in \mathcal{M}$ and $c \in C$. Say

$$Pr[C = c|M = m] = Pr[C = c].$$

Multiplying both sides of the equation by Pr[M = m]/Pr[C = c] gives

$$\frac{Pr[C = c|M = m] \cdot Pr[M = m]}{Pr[C = c]} = Pr[M = m]$$

Using Bayes' theorem, the left-hand-side is exactly equal to $Pr[M = m|C = c]$. Thus, $Pr[M = m|C = c] = Pr[M = m]$ and the scheme is perfectly secret. ∎

Lemma 2 *An encryption scheme (Gen,Enc,Dec) over a message space $\mathcal{M}$ is* perfectly secret *if and only if for every probability distribution over $\mathcal{M}$, every $m_0, m_1 \in \mathcal{M}$, and every cipher text $c \in \mathcal{C}$:*

$$Pr[C = c|M = m_0] = Pr[C = c|M = m_1]$$

Proof: Assume that the encryption scheme is perfectly secret and fix messages $m_0, m_1 \in \mathcal{M}$ and a cipher text $c \in \mathcal{C}$. By Lemma 1 we have

$$Pr[C = c|M = m_0] = Pr[C = c] = Pr[C = c|M = m_1],$$

completing the proof of the first direction.

Assume next that for every distribution over $\mathcal{M}$, every $m_0, m_1 \in \mathcal{M}$, and every $c \in \mathcal{C}$ it holds that $Pr[C = c|M = m_0] = Pr[C = c|M = m_1]$. Fix some distribution over $\mathcal{M}$, and an arbitrary $m_0 \in \mathcal{M}$ and $c \in \mathcal{C}$. Let p be $Pr[C = c|M = m_0]$. Since $Pr[C = c|M = m] = Pr[C = c|M = m_0] = p$ for all m, we have

$$\begin{aligned}
Pr[C = c] &= \sum_{m \in \mathcal{M}} Pr[C = c|M = m] \cdot Pr[M = m] \\
&= \sum_{m \in \mathcal{M}} p \cdot Pr[M = m] \\
&= p \cdot \sum_{m \in \mathcal{M}} Pr[M = m] \\
&= p \\
&= Pr[C = c|M = m_0].
\end{aligned}$$

Since m_0 was arbitrary, we have shown that $Pr[C = c] = Pr[C = c|M = m]$ for all $c \in \mathcal{C}$ and $m \in \mathcal{M}$. Applying Lemma 1, we conclude that the encryption scheme is perfectly secret.

1.2 Shannon's Theorem

Theorem 1 ***(Shannon's Theorem)*** *Let (Gen,Enc,Dec) be an encryption scheme over a message space $\mathcal{M}$ for which $|\mathcal{M}| = |\mathcal{K}| = |\mathcal{C}|$. The scheme is perfectly secret if and only if:*

1. *Every key $k \in \mathcal{K}$ is chosen with equal probability $1/|\mathcal{K}|$ by algorithm Gen.*

2. *For every $m \in \mathcal{M}$ and every $c \in \mathcal{C}$, there exists a unique key $k \in \mathcal{K}$ such that $Enc_k(m)$ outputs c.*

1.3 One-Time Pad Algorithm

Let $a \oplus b$ denote the *bitwise exclusive-or*(XOR) of two binary strings a and b. The one-time pad encryption scheme is as follows:

1. Fix an integer $l > 0$. Then the message space $\mathcal{M}$, key space $\mathcal{K}$, and cipher text space $\mathcal{C}$ are all equal to $\{0, 1\}^l$ (i.e., the set of all binary strings of length l).

2. The key-generation algorithm *Gen* works by choosing a string from $\mathcal{K} = \{0, 1\}^l$ according to the uniform distribution (i.e., each of the 2^l strings in the space chosen as the key with probability exactly 2^{-l}).

3. Encryption *Enc* works as follows: given a key $k \in \{0,1\}^l$ and a message $m \in \{0,1\}^l$, output $c := k \oplus m$.

4. Decryption *Dec* works as follows: given a key $k \in \{0,1\}^l$ and a cipher text $c \in \{0,1\}^l$, output $m := k \oplus c$.

2 Analysis

2.1 Proof of Shannon's Theorem

The intuition behind the proof of Shannon's theorem is as follows. First, if a scheme fulfills item (2) then a given cipher text c could be the result of encrypting any possible plain text m (this holds because for every m there exists a key k mapping it to c). Combining this with the fact that exactly one key maps each m to c, and then by item (1) each key is chosen with the same probability, perfect secrecy can be shown as in the case of the one time pad. For the other direction, the intuition is that if $|\mathcal{M}| = |\mathcal{K}| = |\mathcal{C}|$ then there must be exactly one key mapping each m to each c. (Otherwise, either some m is not mapped to a given c contradicting perfect secrecy, or some m is mapped by more than one key to c, resulting in another m' not being mapped to c, again contradicting perfect secrecy.) Given this, it must hold that each key is chosen with equal probability or some plain texts would be more likely than others, contradicting perfect secrecy. The formal proof follows. Let *(Gen,Enc,Dec)* be as in the theorem. For simplicity, we assume *Enc* is deterministic. We first prove that if *(Gen,Enc,Dec)* is perfectly secret, then items (1) and (2) hold. It is not hard to see that for every $m \in \mathcal{M}$ and $c \in \mathcal{C}$, there exists at least one key $k \in \mathcal{K}$ such that $Enc_k(m) = c$. (Otherwise, $Pr[M = m|C = c] = 0 \neq Pr[M = m]$.) For a fixed m, consider now the set $\{Enc_k(m)\}_{k \in \mathcal{K}}$. By what we have just said, $|\{Enc_k(m)\}_{k \in \mathcal{K}}| \leq |\mathcal{C}|$. We conclude that

$$|\{Enc_k(m)\}_{k \in \mathcal{K}}| = |\mathcal{C}|$$

Since $|\mathcal{K}| = |\mathcal{C}|$, it follows that $\{Enc_k(m)\}_{k \in \mathcal{K}} = |\mathcal{K}|$. This implies that there are no distinct keys $k_1, k_2 \in \mathcal{K}$ with $Enc_{k_1}(m) = Enc_{k_2}(m)$. Since m was arbitrary, we see that for every m and c, there exists at most one key $k \in \mathcal{K}$ such that $Enc_k(m) = c$. Combining the above(i.e., the existence of at least one key and at most one key), we obtain item (2).

We proceed to show that for every $k \in \mathcal{K}$, $Pr[K = k] = 1/\mathcal{K}$. Let $n = |\mathcal{K}|$ and $\mathcal{M} = \{m_1, \ldots, m_n\}$ and fix a cipher text c. Then, we can label the keys $k_1, \ldots, k_2$ so that every $i(1 \leq i \leq n)$ it holds that $Enc_{k_i}(m_i) = c$. This labeling can be carried out because, as just shown, for every c and m_i there exists a unique key k_i such that $Enc_{k_i}(m_i) = c$, and furthermore these keys are distinct for distinct m_i, m_j (since otherwise unambiguous decryption would be impossible). By perfect secrecy we have that for every i:

$$\begin{array}{rcl} Pr[M = m_i] & = & Pr[M = m_i|C = c_i] \\ & = & \dfrac{Pr[C = c_i|M = m_i] \cdot Pr[M = m_i]}{C = c} \\ & = & \dfrac{Pr[K = k_i] \cdot Pr[M = m_i]}{C = c}, \end{array}$$

where the second equality is by Bayes' theorem and the third equality holds by the labeling above (i.e., k_i is the unique key that maps m_i to c). From the above, it follows that for every i,

$$Pr[K = k_i] = Pr[C = c].$$

Therefore, for every i and j, $Pr[K = k_i] = Pr[C = c] = Pr[K = k_j]$ and so all keys are chosen with the same probability. We conclude that keys are chosen according to the uniform distribution. That is, for every k, $Pr[K = k_i] = 1/|\mathcal{K}|$ as required.

We now prove the other direction of the theorem. Assume that every key is obtained with probability $1/\mathcal{K}$ and that for every $m \in \mathcal{M}$ and $c \in \mathcal{C}$ there is a unique key $k \in \mathcal{K}$ such that $Enc_k(m) = c$. This immediately implies that for every m and c,

$$Pr[C = c|M = m_i] = \frac{1}{|\mathcal{K}|}$$

irrespective of the probability distribution over $\mathcal{M}$. Thus, for every probability distribution over $\mathcal{M}$, every $m, m' \in \mathcal{M}$, and every $c \in \mathcal{C}$ we have

$$Pr[C = c|M = m] = \frac{1}{|\mathcal{K}|} = Pr[C = c|M = m'],$$

and so by Lemma 2 the encryption scheme is perfectly secret. ■

2.2 One-Time Pad is Perfectly Secret

Theorem 2 *The one-time pad encryption scheme is perfectly-secret.*

Proof: Fix some distribution over $\mathcal{M}$ and fix an arbitrary $m \in \mathcal{M}$ and $c \in \mathcal{C}$. The key observation is that for the one-time pad,

$$\begin{aligned} Pr[C = c|M = m] &= Pr[M \oplus K = c|M = m] \\ &= Pr[m \oplus K = c] = Pr[K = m \oplus c] = \frac{1}{2^l}. \end{aligned}$$

Since this holds for all distributions and all m, we have that for every probability distribution over $\mathcal{M}$, every $m_0, m_1 \in \mathcal{M}$ and every $c \in \mathcal{C}$,

$$Pr[C = c|M = m_0] = \frac{1}{2^l} = Pr[C = c|M = m_1].$$

By Lemma 2, this implies that the encryption scheme is perfectly secret. ■

3 Impact

3.1 Limitations of One-Time Pad

Unfortunately, the one-time pad encryption scheme has a number of drawbacks. Most prominent is that the key is required to be as long as the message. First and foremost, this means that a long key must be securely stored, something that is highly problematic in practice and often not achievable. In addition, this limits applicability of the scheme if we want to send very long messages (as it may be difficult to securely store a very long key) or if we don't know in advance an upper bound on how long the message will be (since we can't share a key of unbounded length). Moreover, the one-time pad scheme, as the name suggests, is only secure if used once (with the same key). Although we did not yet define a notion of security when multiple messages are encrypted, it is easy to see informally that encrypting more than one message leaks a lot of information. In particular, say two messages m, m' are encrypted using the same key k. An adversary who obtains $c = m \oplus k$ and $c' = m' \oplus k$ can compute

$$\begin{aligned} c \oplus c' &= (m \oplus k) \oplus (m' \oplus k) \\ &= m \oplus m' \end{aligned}$$

and thus learn something about the exclusive-or of the two messages. While this may not seem very significant, it is enough to rule out any claims of perfect secrecy when encrypting two messages. Furthermore, if the messages correspond to English-language text, then given the exclusive-or of two sufficiently long messages, it has been shown to be possible to perform frequency analysis and recover the messages themselves.

4 Implementation

4.1 How to use One-Time Pad encryption in practice

Despite its problems, the one-time-pad retains some practical interest. In some hypothetical espionage situations, the one-time pad might be useful because it can be computed by hand with only pencil and paper. Indeed, nearly all other high quality ciphers (such as Solitaire) are entirely impractical without computers. Spies can receive their pads in person from their "handlers." In the modern world, however, computers (such as those embedded in personal electronic devices such as mobile phones) are so ubiquitous that possessing a computer suitable for performing conventional encryption (for example, a phone which can run concealed cryptographic software) will usually not attract suspicion.

The classical one-time pad of espionage used actual pads of minuscule, easily-concealed paper, a sharp pencil, and some mental arithmetic. The method can be implemented now as a software program, using data files as input (plaintext), output (ciphertext) and key material (the required random sequence). The XOR operation is often used to combine the plaintext and the key elements, and is especially attractive on computers since it is usually a native machine instruction and is therefore very fast. However, ensuring that the key material is actually random, is used only once, never becomes known to the opposition, and is completely destroyed after use is hard to do. The auxiliary parts of a software one-time pad implementation present real challenges: secure handling/transmission of plaintext, truly random keys, and one-time-only use of the key.