Notes for Lectures 1–2

We consider the following scenario: Alice is sending Bob off on an important mission. Prior to Bob's leaving, Alice gives him secret instructions on how to communicate back to her *secretly* so that no eavesdropper Eve can intercept their communication.

The first question to ask is what part of the instructions should we consider secret. To be on the safe side, we must assume that Eve knows as much as possible, and still ensure secrecy under such an assumption.

As far back as 1883, Kerckhoffs [Ker83] suggested that all one can hope to keep secret in a cryptosystem is a key. The algorithms and designs should be assumed to be publicly known. His insight is true to this day, as multiple recent examples demonstrate: publicly known security technologies receive more scrutiny, and hence it is more likely that any problems will be uncovered at early stages. Deploying a secret system and waiting for it to be broken is generally a poor strategy. In this class we will always assume that the adversary knows the entire design of the system.

It is imperative to ask what constitutes a secure system—in other words, what is our goal? Note that, unlike an attack on a system, security cannot be demonstrated by example. Hence, popular perception often holds that "human ingenuity cannot concoct a cipher which human ingenuity cannot resolve" (this quote is from Edgar Allan Poe, who, in addition to being a writer, was an amateur cryptogapher [Poe65]). We will not be satisfied with this design-then-break approach to security. Rather, we will *define* what constitutes a secure system and then *prove* that a particular construction is secure, thus *guaranteeing* security.

The first formal definition of encryption was given by Shannon in his 1949 paper [Sha49].

Definition 1 (encryption scheme a.k.a. cryptosystem). Let M and K be finite sets, and Enc, Dec be two algorithms (Enc may be randomized). We say that $(M, K, \mathrm{Enc}, \mathrm{Dec})$ is an encryption scheme if for all $m \in M$ and $k \in K$, $m = \mathrm{Dec}_k(\mathrm{Enc}_k(m))$ (if Enc is randomized, this equation should hold with probability 1 over the random choices made by Enc).

Note that this definition says nothing about security, it's purely functional. We now address security separately.

Shannon in [Sha49] defined the notion of *perfect security* as follows. Suppose that there is some *a priori* probability distribution D_M on the space of possible messages M (for instance, if the space of messages consists of possible military commands, then some commands are more likely than others). In particular, if the adversary has a guess g for what message is sent, and then Bob picks a message to send, the probability of correctness of the adversary's guess is $\Pr_{m \in D_M}[m = g]$. Now imagine that the adversary first sees a ciphertext c of message chosen by Bob. That defines an *a posteriori* probability distribution on what the message can be, given that its encryption is c. Now, if the adversary has a guess g for what the message is, the probability of the guess being correct, conditioned on c, is $\Pr_{m \in D_m, k \in K}[m = g | \mathrm{Enc}_k(m) = c]$. What the definition says is that the *a priori* and the *a posteriori* probabilities are equal.

Definition 2 (perfect secrecy). An encryption scheme $(M, K, \mathrm{Enc}, \mathrm{Dec})$ satisfies *perfect security with respect to a distribution D_M on M,* if for every message $g \in M$ and for every ciphertext c,

$$\Pr_{m \in D_m, k \in K}[m = g | \mathrm{Enc}_k(m) = c] = \Pr_{m \in D_m}[m = g] \,.$$

(In the above, and from now on, k is chosen *uniformly* from the space of all keys.) An encryption scheme satisfies *perfect security* if for ever D_M it satisfies perfect security with respect to D_M.

We can also consider another notion of security: no matter what message you encrypt, the probability of getting a particular ciphertext is the same. (Intuitively, therefore, the adversary knows nothing about the message from seeing the ciphertext). We will call this definition "Shannon secrecy."

Definition 3 (Shannon secrecy). An encryption scheme $(M, K, \mathrm{Enc}, \mathrm{Dec})$ satisfies *Shannon secrecy* if for every two messages $m_1, m_2 \in M$ and for every ciphertext c,

$$\Pr_{k \in K}[\mathrm{Enc}_k(m_1) = c] = \Pr_{k \in K}[\mathrm{Enc}_k(m_2) = c]\,.$$

Note that both of these definitions consider only the issue of encrypting a *single* message *once*. They do not say what happens if we encrypt two or more messages with the same key, for example.

In [Sha49], Shannon proves that the two definitions are equivalent, which is a good sign: if we have two definitions that both seem intuitively right, then the fact that they are equivalent gives some level of assurance that that our formalizations were correct. (In particular, because these definitions are equivalent, and because Shannon, being modest, never gave the second one a name, you may see the terms "Shannon secrecy" and "perfect secrecy" used interchangeably for the two definitions. In these notes, we use the term "perfect" for the same definition as what Shannon called "perfect" in his paper.)

Theorem 1. *A cryptosystem* $(M, K, \mathrm{Enc}, \mathrm{Dec})$ *satisfies Shannon secrecy if and only if it satisfies perfect secrecy.*

Proof. First, the "only if" direction. Let D_M be a distribution on M, let $g \in M$, and let c be a ciphertext. Then $\Pr_{m,k}[m = g | \mathrm{Enc}_k(m) = c] = \frac{\Pr_{m,k}[\mathrm{Enc}_k(m)=c \wedge m=g]}{\Pr_{m,k}[\mathrm{Enc}_k(m)=c]}$ (by definition of conditional probability). Note that $\Pr_{m,k}[\mathrm{Enc}_k(m) = c \wedge m = g] = \Pr_{m,k}[\mathrm{Enc}_k(g) = c \wedge m = g]$ (we just substituted m for g in the encryption, which we can do, because the condition requires $m = g$). Now, note that the events $\mathrm{Enc}_k(g) = c$ and $m = g$ are independent (because g is fixed, so in the first event the outcome depends entirely on the choice of k, and in the second event the outcome depends entirely on the choice of m). Hence, we get $\Pr_{m,k}[\mathrm{Enc}_k(g) = c \wedge m = g] = \Pr_k[\mathrm{Enc}_k(g) = c]\Pr_m[m = g]$. Finally, note that by Shannon secrecy, the probability that an encryption of g is c is the same as that the encryption of a random message is c: $\Pr_{m,k}[\mathrm{Enc}_k(m) = c] = \sum_{m \in D_M} \Pr[m] \Pr_k[\mathrm{Enc}_k(m) = c] = \sum_{m \in D_m} \Pr[m] \Pr_k[\mathrm{Enc}_k(g) = c] = \Pr_k[\mathrm{Enc}_k(g) = c] \sum_{m \in D_m} \Pr[m] = \Pr_k[\mathrm{Enc}_k(g) = c]$. Putting it all together, we get $\Pr_{m,k}[m = g | \mathrm{Enc}_k(m) = c] = \frac{\Pr_{m,k}[\mathrm{Enc}_k(m)=c]\Pr_m[m=g]}{\Pr_{m,k}[\mathrm{Enc}_k(m)=c]} = \Pr_m[m = g]$, which is perfect secrecy.

Now the "if" direction. Fix any two messages m_1, m_2 and ciphertext c. Pick any distribution D_M that has non-zero probabilities for m_1 and m_2. Then, just like before (except using m_1 for g), we get $\Pr_{m,k}[m = m_1 | \mathrm{Enc}_k(m) = c] = \frac{\Pr_k[\mathrm{Enc}_k(m_1)=c]\Pr_m[m=m_1]}{\Pr_{m,k}[\mathrm{Enc}_k(m)=c]}$, and we know that it's equal, by perfect secrecy, to $\Pr_m[m = m_1]$. Canceling $\Pr_m[m = m_1]$ (it's non-zero because that's how we chose D_M), we get $\frac{\Pr_k[\mathrm{Enc}_k(m_1)=c]}{\Pr_{m,k}[\mathrm{Enc}_k(m)=c]} = 1$. Same for m_2: $\frac{\Pr_k[\mathrm{Enc}_k(m_2)=c]}{\Pr_{m,k}[\mathrm{Enc}_k(m)=c]} = 1$. Because the fractions are equal for m_1 and m_2 and the denominators are the same, the numerators must be equal as well: $\Pr_k[\mathrm{Enc}_k(m_1) = c] = \Pr_k[\mathrm{Enc}_k(m_2) = c]$. $\square$

Consider now the following scheme that satisfies Shannon secrecy for $M = \{0, 1\}$: the key space is the space of the two permutations of $\{0, 1\}$, and encryption is defined as $\mathrm{Enc}_k(m) = k \oplus m$. The proof that it satisfies Shannon secrecy is quite simple, because there are only two messages and two ciphertexts: $\Pr_k[\mathrm{Enc}_k(0) = 0] = \Pr_k[\mathrm{Enc}_k(1) = 0] = \Pr_k[\mathrm{Enc}_k(0) = 1] = \Pr_k[\mathrm{Enc}_k(1) = 1] = 1/2$.

We now generalize this scheme for longer messages by repeating the scheme for one-bit messages. The result is known as the one-time-pad, or the Vernam cipher (patented by Vernam in 1919 [Ver19] and published in 1926 [Ver26], for use with punch-tape in telegraph transmission). For any integer n, let $M = K = \{0,1\}^n$, and let $\mathrm{Enc}_k(m) = m \oplus k$. Let $\mathrm{Dec}_k(c) = c \oplus k$. It's easy to see that for any fixed $m \in M$ and $c \in \{0,1\}^n$, $\mathrm{Pr}_{k\in K}[\mathrm{Enc}_k(m) = c] = Pr_{k\in K}[k = m \oplus c] = 2^{-n}$. Hence, for any $m_1 \in M, m_2 \in M, c$, $\mathrm{Pr}_k[\mathrm{Enc}_k(m_1) = c] = \mathrm{Pr}_k[\mathrm{Enc}_k(m_2) = c]$.

This is a very computationally efficient scheme: exclusive-or is a simple operation, and multiple bits can be encrypted independently in parallel. Its main drawback is that it requires a very long key, which can be used only once. Next time we need to encrypt a message, we need to select a new random key (otherwise, the adversary could, for example, compute the exclusive-or of the two messages encrypted).

Shannon showed that, unfortunately, it's the best one can do. The following theorem is known as the Shannon bound.

Theorem 2. *If* $(M, K, \mathrm{Enc}, \mathrm{Dec})$ *satisfies perfect secrecy, then* $|K| \geq |M|$.

Proof. Let c be a possible ciphertext—i.e., fix some message $m_1 \in M$, and let c be such that $\Pr[\mathrm{Enc}_{k_1}(m_1) = c] > 0$. Suppose there is some $m_2 \in M$ such that for all $k \in K$, $\mathrm{Dec}_k(c) \neq m_2$. Then, by definition of encryption scheme m_2 would never get encrypted to c (because otherwise you couldn't decrypt it). So $\Pr[\mathrm{Enc}_{k_1}(m_2) = c] = 0$, so $\Pr[\mathrm{Enc}_{k_1}(m_1) = c] \neq \Pr[\mathrm{Enc}_{k_1}(m_2) = c]$, which violates perfect secrecy. In other words, c must be decryptable to all plaintexts in M. Hence, for each $m_2 \in M$, there exists $k \in K$ such that $\mathrm{Dec}_k(c) = m_2$. So there must be at least as many $k \in K$ as $m_2 \in M$, so $|K| \geq |M|$. □

This pretty much ends our discussion of information-theoretic cryptography: we have an efficient encryption scheme, and a proof that you can't do better. We will have to limit adversary's computational power if we want anything more efficient.

References

[Ker83] Auguste Kerckhoffs. La cryptographie militaire. *Journal des Sciences Militaires*, 9th Series:5–38 and 161–191, January-February 1883. Available at `http://www.petitcolas.net/fabien/kerckhoffs/`.

[Poe65] Edgar Allan Poe. A few words on secret writing. In James A. Harrison, editor, *The Complete Works of Edgar Allan Poe*, volume XIV, pages 114–149. AMS Press, NY, 1965. Originally in Graham's Magazine, July 1841.

[Sha49] Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, October 1949. Available at `http://www.cs.ucla.edu/~jkong/research/security/shannon.html` and `http://www3.edgenet.net/dcowley/docs.html`.

[Ver19] Gilbert S. Vernam. U.S.Patent 1,310,719. Secret signaling system, 22 July 1919.

[Ver26] Gilbert S. Vernam. Cipher printing telegraph systems for secret wire and radio telegraphic communications. *Journal of the American Institute of Electrical Engineers*, 45:109–115, 1926.