Adversarial indistinguishability
Computationally-secure private-key encryption

Foundations of Cryptography
Computer Science Department
Wellesley College

Fall 2016

Table of contents

The syntax of encryption

Definition 3.7. A *private-key encryption scheme* is a tuple of probabilistic polynomial-time algorithms (Gen, Enc, Dec) such that

1. The *key-generation algorithm* Gen takes as input the security parameter 1^n and outputs a key k; we write $k \leftarrow \mathrm{Gen}(1^n)$ and assume WLOG that any key k output by $\mathrm{Gen}(1^n)$ satisfies $|k| \geq n$.
2. The *encryption algorithm* Enc takes as input a key k and a plaintext message $m \in \{0,1\}^*$ and outputs a ciphertext c. We write $c \leftarrow \mathrm{Enc}_k(m)$.
3. The *decryption algorithm* Dec takes as input a key k and a ciphertext c and output a plaintext m. We write $m := \mathrm{Dec}_k(c)$.

The eavesdropping adversary

Any definition of security consists of two distinct components:

1. A specification of the assumed power of the adversary;
2. And a description of what constitutes a "break";

We begin by considering the case of an *eavesdropping adversary* who observes the encryption of a single message.

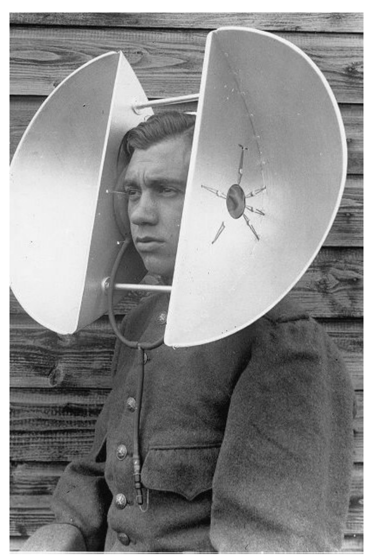

Never underestimate your adversary

- Although we assume our adversary only eavesdrops and runs in polynomial time, we make no assumptions about the adversary's strategy.
- Since we cannot predict all possible strategies, we must protect against any possible attack within the class defined.

What's a break?

- Defining the "break" isn't easy, however, we already agreed that the adversary should be unable to learn *any partial information* about the plaintext from the ciphertext.
- The definition of *semantic security* formalizes this notion, but is difficult to work with.
- Fortunately, there is an equivalent definition using *indistinguishability* which is much simpler.

More experiments in security

The experiment is defined for any private-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$, any adversary $\mathcal{A}$, and any value n for the security parameter:

The eavesdropping indistinguishability experiment $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n)$

1. The adversary $\mathcal{A}$ is given 1^n, and outputs a pair of messages $m_0, m_1 \in \mathcal{M}$ of the same length.
2. A key k is generated by running $\mathsf{Gen}(1^n)$, and a random bit $b \leftarrow \{0, 1\}$ is chosen. A *challenge ciphertext* $c \leftarrow \mathsf{Enc}_k(m_b)$ is computed and given to $\mathcal{A}$.
3. $\mathcal{A}$ outputs a bit b'.
4. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise. We write $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n) = 1$ if the output is 1 and in this case we say that $\mathcal{A}$ *succeeded*.

Adversarial indistinguishability

Definition 3.8. An encryption scheme Π = (Gen, Enc, Dec) had *indistinguishable encryption in the presence of an eavesdropper* if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there exists a negligible function negl such that

$$\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n) = 1] \leq \frac{1}{2} + \mathsf{negl}(n),$$

where the probability is taken over the random coins used by $\mathcal{A}$, as well as the random coins used by the experiment (for choosing the key, the random bit b, and any random coins used in the encryption process).

Put another way

- Definition 3.8 states that an eavesdropping adversary cannot determine which plaintext was encrypted with probability better than guessing.
- Another way to say this is that every adversary behaves the same way whether it sees an encryption of m_0 or an encryption of m_1
- We formalize this notion in the following definition.

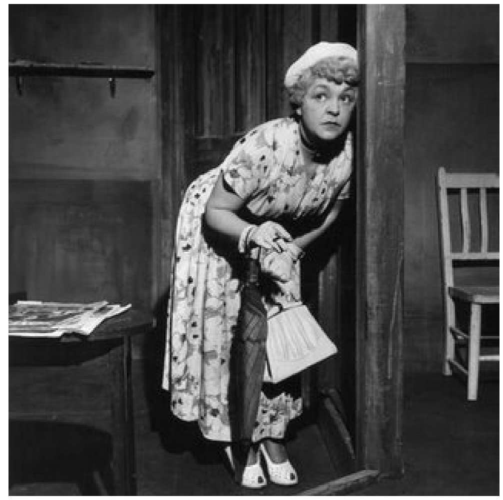

Indistinguishable encryptions again

Define $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n, b)$ to be as above, except the fixed bit b is used. In addition denote the output bit b' of $\mathcal{A}$ in $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n, b)$ by $\mathsf{output}(\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n, b))$.

Definition 3.9. An encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ had *indistinguishable encryption in the presence of an eavesdropper* if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there exisits a negligible function negl such that

$$\left|\Pr[\mathsf{output}(\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n, 0) = 1] - \Pr[\mathsf{output}(\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n, 1) = 1]\right| \leq \mathsf{negl}(n)$$

Bait and switch

- We motivated the definition of secure encryption by saying that it should be infeasible for an adversary to learn any partial information about the plaintext from the ciphertext (*semantic security*).
- But our definition doesn't look anything like that.
- We prove two claims demonstrating our definition isn't so far off the mark.

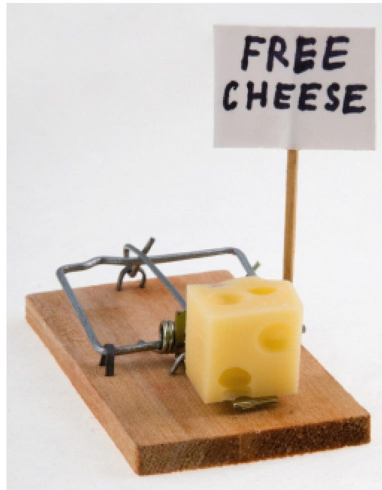

No random bit of the plaintext can be determined better than guessing

Denote by m^i the ith bit of m, and set $m^i = 0$ if $i > |m|$.

Claim 3.10. Let (Gen, Enc, Dec) be a private-key encryption scheme that has indistinguishable encryption in the presence of an eavesdropper. Then for all probabilistic polynomial-time adversaries $\mathcal{A}$ and all i, there exists a negligible function negl such that:

$$\Pr[\mathcal{A}(1^n, \mathsf{Enc}_k(m)) = m^i] \le \frac{1}{2} + \mathsf{negl}(n)$$

where m is chosen uniformly at random from $\{0,1\}^n$, and the probability is taken over the random coins of $\mathcal{A}$, the choice of m and the key k, and any random coins used in the encryption process.

Proving Claim 3.10

Proof. Let $\mathcal{A}$ be a probabilistic polynomial-time adversary and define $\epsilon(\cdot)$ as follows:

$$\epsilon(n) \stackrel{\text{def}}{=} \Pr[\mathcal{A}(1^n, \mathsf{Enc}_k(m)) = m^i] - \frac{1}{2}.$$

where m is chosen uniformly from $\{0,1\}^n$.
Take $n \geq i$, let I_0^n be the set of all strings of length n whose ith bit is 0. Likewise I_1^n. It follows that:

$$\Pr[\mathcal{A}(\mathsf{Enc}_k(m)) = m^i] = \frac{1}{2}\cdot\Pr[\mathcal{A}(\mathsf{Enc}_k(m_0)) = 0] + \frac{1}{2}\cdot\Pr[\mathcal{A}(\mathsf{Enc}_k(m_1)) = 1]$$

where m_0, m_1 are chose uniformly from I_0^n, I_1^n respectively.

Consider the following eavesdropping adversary $\mathcal{A}'$

Adversary $\mathcal{A}'$:

1. On input 1^n (with $n \geq i$), choose $m_0 \leftarrow I_0^n$ and $m_1 \leftarrow I_1^n$ uniformly and output m_0, m_1.
2. Upon receiving a ciphert text c, invoke $\mathcal{A}$ on input c. Output $b' = 0$ if $\mathcal{A}$ outputs 0, and $b' = 1$ if $\mathcal{A}$ outputs 1.

$\mathcal{A}'$ runs in polynomial time since $\mathcal{A}$ does. Using the definition of $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A}',\Pi}(n)$, note that $b' = b$ if and only if $\mathcal{A}$ outputs b upon receiving $\mathsf{Enc}_k(m_b)$. So

$$\begin{aligned}\Pr[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A}',\Pi}(n) = 1] &= \Pr[\mathcal{A}(\mathsf{Enc}_k(m_b)) = b] \\ &= \frac{1}{2}\cdot\Pr[\mathcal{A}(\mathsf{Enc}_k(m_0)) = 0] + \frac{1}{2}\cdot\Pr[\mathcal{A}(\mathsf{Enc}_k(m_1)) = 1] \\ &= \Pr[\mathcal{A}(\mathsf{Enc}_k(m)) = m^i] = \frac{1}{2} + \epsilon(n).\end{aligned}$$

Since (Gen, Enc, Dec) has indistinguishable encryptions in the presence of an eavesdropper, $\epsilon(\cdot)$ must be negligible. □

Finding the correct definition of semantic security

- We wish that no PPT adversary can learn *any* function of the plaintext given the ciphertext regardless of the *a priori* distribution of messages sent.
- But even computing the ith bit of the plaintext m is easy when m is chosen uniformly from I_0^n.
- What we want to say is that an adversary receiving $c = \mathsf{Enc}_k(m)$ can compute $f(m)$, there there exists an adversary that can compute $f(m)$ with the same probability without being given c.

Close to semantic security

Claim 3.11. Let (Gen, Enc, Dec) be a private-key encryption scheme that has indistinguishable encryption in the presence of an eavesdropper. Then for every PPT adversary $\mathcal{A}$ there exists a PPT adversary $\mathcal{A}'$ such that for all polynomial-time computable functions f and all efficiently-sampleable sets S, there exists a negligible function negl such that:

$$\left|\Pr[\mathcal{A}(1^n, \mathsf{Enc}_k(m)) = f(m)] - \Pr[\mathcal{A}'(1^n) = f(m)]\right| \leq \mathsf{negl}(n)$$

where m is chosen uniformly at random from $S_n \stackrel{\text{def}}{=} S \cap \{0,1\}^n$, and the probability is taken over the random coins of $\mathcal{A}$, the choice of m and the key k, and any random coins used by the adversaries and encryption process.

*Since we are considering an asymptotic setting, we work with an infinite set $S \subseteq \{0.1\}^*$.

Sketch of proof of Claim 3.11

- Suppose (Gen, Enc, Dec) has indistinguishable encryption in the presence of an eavesdropper. Then for no PPT adversary $\mathcal{A}$ can distinguish between $\mathsf{Enc}_k(m)$ and $\mathsf{Enc}_k(1^n)$ for any $m \in \{0,1\}^n$.
- Consider the probability that $\mathcal{A}$ successfully computes $f(m)$ given $\mathsf{Enc}_k(m)$. $\mathcal{A}$ should successfully compute $f(m)$ given $\mathsf{Enc}_k(1^n)$ with almost the same probability. Otherwise, $\mathcal{A}$ could be used to distinguish between $\mathsf{Enc}_k(m)$ and $\mathsf{Enc}_k(1^n)$
- Construct algorithm $\mathcal{A}'$: On input 1^n, choose a random key k, invoke $\mathcal{A}$ on $c \leftarrow \mathsf{Enc}_k(1^n)$, and output whatever $\mathcal{A}$ does. By above, subroutine $\mathcal{A}$ outputs $f(m)$ with same probability as when it receives $\mathsf{Enc}_k(m)$. $\square$