Public-key encryption
The details

Foundations of Cryptography
Computer Science Department
Wellesley College

Fall 2016

Table of contents

Public-key encryption scheme

Definition 11.1. A *public-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 pair of keys (pk, sk) with $|pk| = n = |sk|$. We refer to these as the *public key* and the *private key* respectively.
2. The *encryption algorithm* Enc takes as input a public key pk and a message m from some underlying plaintext space. It outputs a ciphertext c; we write $c \leftarrow \mathsf{Enc}_{pk}(m)$.
3. The *decryption algorithm* Dec*:* takes as input a private key sk and a ciphertext c, and outputs a message m or a special symbol $\perp$ denoting failure. We assume WLOG that Dec is deterministic and write $m := \mathsf{Dec}_{sk}(c)$.

We require that, except with negligible probability,

$$\mathsf{Dec}_{sk}(\mathsf{Enc}_{pk}(m)) = m$$

The eavesdropping indistinguishability experiment

Given a public-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ and an adversary $\mathcal{A}$ consider the following:

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

1. $\mathsf{Gen}(1^n)$ is run to obtain keys (pk, sk).
2. Adversary $\mathcal{A}$ is given pk, and outputs a pair of messages m_0, m_1 of the same length.
3. A random bit $b \leftarrow \{0, 1\}$ is chosen, and then a ciphertext $c \leftarrow \mathsf{Enc}_{pk}(m_b)$ is computed and given to $\mathcal{A}$. We call c the *challenge ciphertext*.
4. $\mathcal{A}$ outputs a bit b'.
5. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise.

*Giving pk to $\mathcal{A}$ effectively gives $\mathcal{A}$ encryption oracle access for free.

Indistinguishable encryptions in the presence of an eavesdropper

Definition 11.2. A public-key encryption scheme $\mathsf{PubK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n)$ has *indistinguishable encryptions in the presence of an eavesdropper* if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there exists a negligible function negl such that

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

Storming the Bastille

- Of course there is more than one form of attack ...
- And hence, more one definition of security.
- For example, we may wish our public-key encryption schemes for be secure against CPA or even CCA attacks.

More experiments and definitions

The CPA indistinguishability experiment $\mathsf{PubK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n)$:

1. $\mathsf{Gen}(1^n)$ is run to obtain keys (pk, sk).
2. Adversary $\mathcal{A}$ is given pk as well as oracle access to $\mathsf{Enc}_{pk}(\cdot)$. The adversary outputs a pair of messages m_0, m_1 of the same length.
3. A random bit $b \leftarrow \{0, 1\}$ is chosen, and then a ciphertext $c \leftarrow \mathsf{Enc}_{pk}(m_b)$ is computed and given to $\mathcal{A}$.
4. $\mathcal{A}$ continues to have access to $\mathsf{Enc}_{pk}(\cdot)$, and outputs a bit b'.
5. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise.

Definition. A public-key encryption scheme $\mathsf{PubK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n)$ has *indistinguishable encryptions under a chosen-plaintext attack* if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there exists a negligible function negl such that

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

But you told me ...

Proposition 11.3 If a public-key encryption scheme $\mathsf{PubK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n)$ has indistinguishable encryptions in the presence of an eavesdropper then Π also has indistinguishable encryptions under a chosen plain-text attack.

Perfectly-secret public-key encryption

Definition. A public-key encryption scheme $\mathsf{PubK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n)$ is *perfectly secret* if for every PPT adversary $\mathcal{A}$

$$\Pr[\mathsf{PubK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n) = 1] = \frac{1}{2}.$$

Sad but true. Unfortunately, perfectly-secret public-key encryption schemes are pipe dreams.*

*We leave this for an exercise.

More pipes: Insecurity of deterministic public-key encryption

Remark. For the same reason that no deterministic private-key encryption scheme can be CPA-secure, we have

Theorem 11.7. No deterministic public-key encryption scheme has indistinguishability in the presence of an eavesdropper

Warning! This is not a mere "artifact" our security definition. Deterministic public-key encryption schemes are vulnerable to *practical* attacks in *realistic* scenarios.

CPA security for multiple encryptions

The definition for indistinguishable encryptions under a chosen-plaintext can easily be extended to indistinguishable multiple encryptions in the same way that indistinguishability encryption in the presence of an eavesdropper was.

The text takes a somewhat simpler approach that can model attackers that can adaptively choose plaintexts to be encrypted, even after observing previous ciphertext.

The attacker has access to a *"left-or-right" oracle $LR_{k,b}$* that, on input a pair of equal-length messages m_0, m_1, computes the ciphertext $c \leftarrow \mathsf{Enc}_k(m_b)$ and returns c.*

*Here b is a random bit chosen at the beginning of the experiment.

One more experiment

The LR-oracle experiment $\mathsf{PubK}^{LR\text{-}cpa}_{\mathcal{A},\Pi}(n)$:

1. $\mathsf{Gen}(1^n)$ is run to obtain keys (pk, sk).
2. A uniform bit $b \leftarrow \{0,1\}$ is chosen.
3. The adversary $\mathcal{A}$ is given input pk and oracle access to $\mathsf{LR}_{pk,b}(\cdot,\cdot)$.
4. Adversary $\mathcal{A}$ outputs a bit b'.
5. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise. If $\mathsf{PubK}^{\mathsf{LR\text{-}cpa}}_{\mathcal{A},\Pi}(n) = 1$, we say that $\mathcal{A}$ succeeds.

Definition 11.5. A public-key encryption scheme $\mathsf{PubK}^{\mathsf{PR\text{-}cpa}}_{\mathcal{A},\Pi}(n)$ has *indistinguishable multiple encryptions* if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there exists a negligible function negl such that

$$\Pr[\mathsf{PubK}^{\mathsf{LP\text{-}cpa}}_{\mathcal{A},\Pi}(n) = 1] \leq \frac{1}{2} + \mathsf{negl}(n).$$

CPA-secure implies indistinguishable multiple encryptions

Theorem 11.6. If a public-key encryption scheme Π is CPA-secure then Π has indistinguishable multiple encryptions.

Remark. Theorem 11.6 implies that a CPA-secure public-key encryption scheme for *fixed-length* messages implies a public-key encryption scheme for *arbitrary-length* messages satisfying the same notion of security.

Remark. For example, suppose $\Pi = (\text{Gen}, \text{Enc}, \text{Dec})$ is an encryption scheme for a single-bit message. We construction $\Pi' = (\text{Gen}, \text{Enc'}, \text{Dec'})$ for messages in $\{0,1\}^*$

$$\text{Enc}'_{pk}(m) = \text{Enc}_{pk}(m_1), \ldots, \text{Enc}_{pk}(m_\ell),$$

where $m = m_1, \ldots, m_\ell$.

Intuition behind Theorem 11.6

Theorem 11.6. If a public-key encryption scheme Π is CPA-secure then Π has indistinguishable multiple encryptions.

Proof. Fix an arbitrary PPT adversary $\mathcal{A}$ and a CPA-secure public-key encryption scheme Π. Consider experiment $\text{PubK}^{\text{LR-cpa2}}_{\mathcal{A},\Pi}(n)$ where $\mathcal{A}$ can only make two queries: $(m_{1,0}, m_{1,1})$ and $(m_{2,0}, m_{2,1})$. In the experiment $\mathcal{A}$ receives either the pair

$$(\text{Enc}_{pk}(m_{1,0}), \text{Enc}_{pk}(m_{2,0})) \text{ or } (\text{Enc}_{pk}(m_{1,1}), \text{Enc}_{pk}(m_{2,1})).$$

We write $\mathcal{A}(pk, \text{Enc}_{pk}(m_{1,0}), \text{Enc}_{pk}(m_{2,0}))$ in the first case and analogously for the second.
We show that there exists a negligible function negl such that

$$\begin{aligned} |\Pr[&\mathcal{A}(pk, \text{Enc}_{pk}(m_{1,0}), \text{Enc}_{pk}(m_{2,0})) = 1] \\ &- \Pr[\mathcal{A}(pk, \text{Enc}_{pk}(m_{1,1}), \text{Enc}_{pk}(m_{2,1})) = 1]| \leq \text{negl}(n). \end{aligned}$$

*For simplicity we assume the adversary make only two calls to the LR oracle.

To prove this, we will show that

Let $\vec{C}_0$ denote the distribution of ciphertext pairs $(\mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,0}))$, and $\vec{C}_1$ the distribution of ciphertext pairs $(\mathsf{Enc}_{pk}(m_{1,1}), \mathsf{Enc}_{pk}(m_{2,1}))$. We show

1. CPA-security of Π implies $\mathcal{A}$ cannot distinguish between when it is give a pair of ciphertexts distributed according to $\vec{C}_0$, or a pair of ciphertext $(\mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,1}))$. Denote the distribution of these ciphertexts by $\vec{C}_{01}$.
2. Similarly, CPA-security of Π implies that $\mathcal{A}$ cannot distinguish between when it is give a pair of ciphertexts distributed according to $\vec{C}_{01}$, or a pair distributed according to $\vec{C}_1$.

We conclude that $\mathcal{A}$ cannot distinguish between distributions $\vec{C}_0$ and $\vec{C}_1$.

The long and the short of it

We must show that there is a negligible function negl for which

$$\begin{aligned} &|\Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,0})) = 1] \\ &\quad - \Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,1})) = 1]| \le \mathsf{negl}(n).* \end{aligned}$$

*Intuitively this follows from the single message case since these two inputs differ only in the second element and $\mathcal{A}$ can generate $\mathsf{Enc}_{pk}(m_{1,0})$ on its own.

To prove our claim, consider the following PPT adversary

Adversary $\mathcal{A}'$ against the single message experiment $PubK^{eav}_{\mathcal{A}',\Pi}(n)$:

1. $\mathcal{A}'$, given pk, runs $\mathcal{A}(pk)$.
2. When $\mathcal{A}(pk)$ makes its first query $(m_{1,0}, m_{1,1})$ to the LR oracle, $\mathcal{A}'$ computes $c_1 \leftarrow \mathsf{Enc}_{pk}(m_{1,0})$ and returns c_1 to $\mathcal{A}$.
3. When $\mathcal{A}(pk)$ makes its second query $(m_{2,0}, m_{2,1})$ to the LR oracle, $\mathcal{A}'$ outputs $(m_{2,0}, m_{2,1})$ and receives back a challenge ciphertext c_2. This is returned to $\mathcal{A}$.
4. $\mathcal{A}'$ outputs the bit b' that is output by $\mathcal{A}$.

When $b = 0$ adversary $\mathcal{A}'$ is given $\mathsf{Enc}_{pk}(m_{2,0})$, and

$$\Pr[\mathcal{A}'(\mathsf{Enc}_{pk}(m_{2,0})) = 0] = \Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,0})) = 0].$$

When $b = 1$ adversary $\mathcal{A}'$ is given $\mathsf{Enc}_{pk}(m_1^2)$, and

$$\Pr[\mathcal{A}'(\mathsf{Enc}_{pk}(m_{2,1})) = 1] = \Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,1})) = 1].$$

Completing the proof of Claim 10.8

By the security of Π in the sense of single-message indistinguishability, there exists a negligible function negl such that

$$\begin{aligned}\frac{1}{2} + \mathsf{negl}(n) &\geq \Pr[\mathsf{PubK}^{eav}_{\mathcal{A}',\Pi}(n) = 1] \\ &= \frac{1}{2} \cdot (\Pr[\mathcal{A}'(\mathsf{Enc}_{pk}(m_{2,0})) = 0] + \Pr[\mathcal{A}'(\mathsf{Enc}_{pk}(m_{2,1})) = 1]) \\ &= \frac{1}{2} \cdot (1 - \Pr[\mathcal{A}'(\mathsf{Enc}_{pk}(m_{2,0})) = 1] + \Pr[\mathcal{A}'(\mathsf{Enc}_{pk}(m_{2,1})) = 1])\end{aligned}$$

So that (after some fiddling to get other side of the absolute values)

$$|\Pr[\mathcal{A}'(\mathsf{Enc}_{pk}(m_{2,0})) = 1] - \Pr[\mathcal{A}'(\mathsf{Enc}_{pk}(m_{2,1})) = 1]| \leq \mathsf{negl}(n).$$

This, together with results from previous page proves:

$$\begin{aligned}&|\Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,0})) = 1] \\ &\quad - \Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,1})) = 1]| \leq \mathsf{negl}(n).\end{aligned}$$

A very similar argument proves

From the previous slide we have

$$
\begin{aligned}
&|\Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,0})) = 1] \\
&\quad - \Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,1})) = 1]| \le \mathsf{negl}(n).
\end{aligned}
$$

An almost identical arguments show the existence of a negligible function negl such that

$$
\begin{aligned}
&|\Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,1})) = 1] \\
&\quad - \Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,1}), \mathsf{Enc}_{pk}(m_{2,1})) = 1]| \le \mathsf{negl}(n).
\end{aligned}
$$

Combining these two inequalities yields

$$
\begin{aligned}
&|\Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,0}), \mathsf{Enc}_{pk}(m_{2,0})) = 1] \\
&\quad - \Pr[\mathcal{A}(pk, \mathsf{Enc}_{pk}(m_{1,1}), \mathsf{Enc}_{pk}(m_{2,1})) = 1]| \le \mathsf{negl}(n).
\end{aligned}
$$

$\square$