Lecture 7: 2/12/02

1 Review

In the last couple of lectures, we have discussed relations between different notions of security. Today we will show one final relation and then introduce symmetric key cryptography. Recall the definitions of security under message non-malleability and message indistinguishability.

Definition 1 *A public-key encryption scheme $\mathcal{PE} = (G, E, D)$ is (t,q,ϵ) secure in the sense of IND-ATK if for all pairs of messages (m_0, m_1), $m_0 \neq m_1$, and for every adversary A running in time t, making at most q queries to oracle O,*

$$\Pr_{(pk,sk)\in G}[A^O(pk, E_{pk}(m_0)) = 1] - \Pr_{(pk,sk)\in G}[A^O(pk, E_{pk}(m_1)) = 1] \leq \epsilon$$

where the oracle O is

$$O = \begin{cases} D_{sk} & \text{if ATK is CCA2} \\ - & \text{if ATK is CPA} \end{cases}$$

and A is restricted from querying the oracle O on its input $E_{pk}(m_i)$.

Definition 2 *A public-key encryption scheme $\mathcal{PE} = (G, E, D)$ is (t,q,ϵ) secure in the sense of NM-ATK if for all message distributions $\mathcal{M}$, and for all relations $\mathcal{R} : \mathcal{M} \times \mathcal{M} \to \{0, 1\}$, and for every adversary A running in time t, making at most q queries to oracle O, there exists a simulator S running in time $poly(t)$, such that*

$$\Pr_{\substack{(pk,sk)\in G \\ m\in\mathcal{M}}}[\mathcal{R}(m, D_{sk}(A^O(pk, E_{pk}(m))))] - \Pr_{\substack{(pk,sk)\in G \\ m\in\mathcal{M}}}[\mathcal{R}(m, D_{sk}(S(pk)))] \leq \epsilon$$

where the oracle O is

$$O = \begin{cases} D_{sk} & \text{if ATK is CCA2} \\ - & \text{if ATK is CPA} \end{cases}$$

and A is restricted from querying the oracle O on its input $E_{pk}(m)$.

2 Equivalence of NM-CCA2 and IND-CCA2

Two lectures ago we showed that NM-CCA2 $\Rightarrow$ IND-CCA2. Today we will complete the equivalence by showing IND-CCA2 $\Rightarrow$ NM-CCA2.

Theorem 1 (IND-CCA2 $\Rightarrow$ NM-CCA2)
If public key encryption scheme $\mathcal{PE} = (G, E, D)$ is (t, q, ϵ) secure in the sense of IND-CCA2, then $\mathcal{PE}$ is $(t, q, 2\epsilon)$ secure in the sense of NM-CCA2.

Proof:
The proof will proceed by contradiction. Suppose $\mathcal{PE}$ is not secure in the sense of NM-CCA2, i.e. there exists a message distribution $\mathcal{M}$, a relation $\mathcal{R} : \mathcal{M} \times \mathcal{M} \to \{0, 1\}$, and a $(t, q, 2\epsilon)$ NM-CCA2 adversary A with running time t, such that for all simulators S running in time at most $poly(t)$,

$$\Pr_{\substack{(pk,sk)\in G\\ m\in\mathcal{M}}}[\mathcal{R}(m, D_{sk}(A^{D_{sk}}(pk, E_{pk}(m))))] - \Pr_{\substack{(pk,sk)\in G\\ m\in\mathcal{M}}}[\mathcal{R}(m, D_{sk}(S(pk)))] > 2\epsilon$$

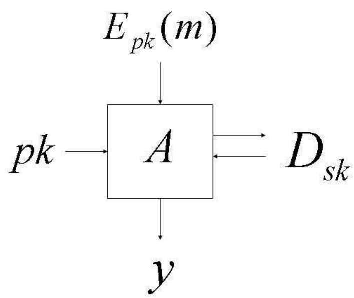

The adversary A and a simulator S can be visualized as above, where both A and S try to output a y such that $\mathcal{R}(m, D_{sk}(y))$.

Let m_0 and m_1 be two messages chosen at random from $\mathcal{M}$ such that $m_0 \neq m_1$, and let B be the following adversary:

Adversary $B(pk, E_{pk}(m))$**, where** $m \in \{m_0, m_1\}$
    Run $A^{D_{sk}}(pk, E_{pk}(m)) \to y$
    If $\mathcal{R}(m_0, D_{sk}(y))$
        return 0
    else
        return $r \leftarrow_R \{0, 1\}$

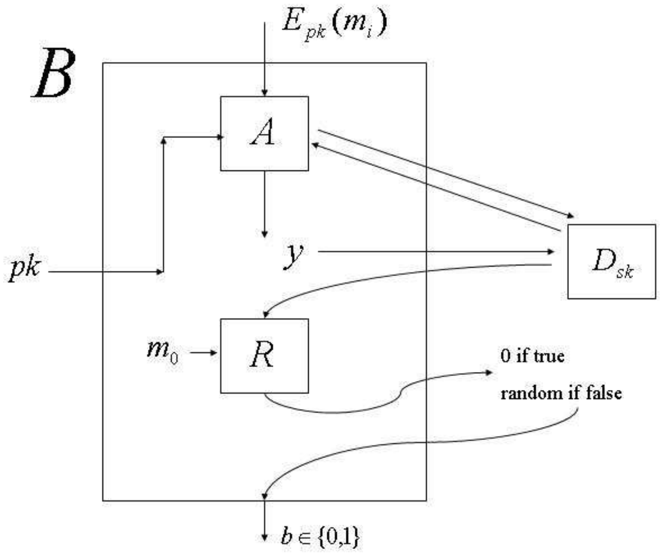

We would like B to be restricted from using the decryption oracle on $E_{pk}(m)$. In B, this would only happen if A output $E_{pk}(m)$, in which case $\mathcal{R}$ would be the trivial equality relation, and this can assumed to be not the case. Define the following simulator S as follows:

Simulator $S(pk)$
Generate $m'' \leftarrow_R \mathcal{M}$
Return $y = A^{D_{sk}}(pk, E_{pk}(m''))$

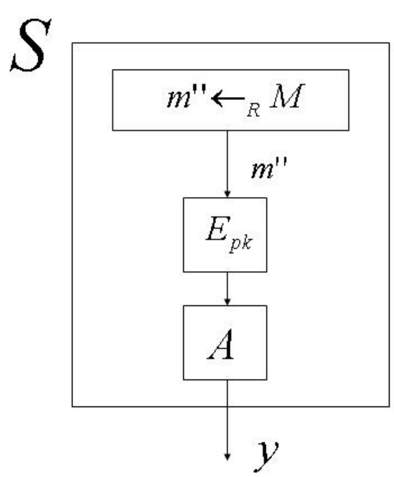

Define

$$p = \Pr_{\substack{(pk,sk)\in G \\ m_0,m_1 \leftarrow \mathcal{M}}} [\mathcal{R}(m_0, D_{sk}(A^{D_{sk}}(pk, E_{pk}(m_0))))]$$

$$p' = \Pr_{\substack{(pk,sk)\in G \\ m_0,m_1 \leftarrow \mathcal{M}}} [\mathcal{R}(m_1, D_{sk}(A^{D_{sk}}(pk, E_{pk}(m_0))))]$$

Note

$$\begin{aligned} p' &= \Pr_{\substack{(pk,sk)\in G \\ m_0,m_1 \leftarrow \mathcal{M}}} [\mathcal{R}(m_0, D_{sk}(A^{D_{sk}}(pk, E_{pk}(m_1))))] \\ &= \Pr_{\substack{(pk,sk)\in G \\ m',m'' \leftarrow \mathcal{M}}} [\mathcal{R}(m', D_{sk}(A^{D_{sk}}(pk, E_{pk}(m''))))] \\ &= \Pr_{\substack{(pk,sk)\in G \\ m' \leftarrow \mathcal{M}}} [\mathcal{R}(m', D_{sk}(S(pk)))] \end{aligned}$$

Since we know that A is a good NM adversary, then $p - p' > 2\epsilon$. The advantage of B as an indistinguishability adversary, for a pair of messages (m_0, m_1), is given by

$$\mathbf{Adv}^{\text{IND-CCA2}}_{B,\mathcal{PE}} = \Pr_{(pk,sk)\in G}[B(E(pk, E_{pk}(m_0))) = 0] - \Pr_{(pk,sk)\in G}[B(E(pk, E_{pk}(m_1))) = 0]$$

Note for $i \in \{0,1\}$, $c_i \leftarrow E_{pk}(m_i)$,

$$\Pr_{(pk,sk)\in G}[B(pk, c_i) = 0] = \Pr_{(pk,sk)\in G}[\mathcal{R}(m_0, D_{sk}(A(pk, c_i)))] + \frac{1}{2}\Pr_{(pk,sk)\in G}[\text{not } \mathcal{R}(m_0, D_{sk}(A(pk, c_i)))]$$

and

$$\Pr_{\substack{(pk,sk)\in G \\ m_0,m_1 \leftarrow \mathcal{M}}} [B(pk, E_{pk}(m_0)) = 0] = p + \frac{1}{2}(1-p)$$

$$\Pr_{\substack{(pk,sk)\in G \\ m_0,m_1 \leftarrow \mathcal{M}}} [B(pk, E_{pk}(m_1)) = 0] = p' + \frac{1}{2}(1-p')$$

Therefore, the average advantage of B over all possible choices of m_0 and m_1 is

$$\begin{aligned} \mathbf{Avg\ Adv}^{\text{IND-CCA2}}_{B,\mathcal{PE}} &= (p + \frac{1}{2}(1-p)) - (p' + \frac{1}{2}(1-p')) \\ &= \frac{1}{2}(p - p') \\ &> \epsilon \end{aligned}$$

This is a contradiction, since by assumption, $\mathcal{PE}$ is (t, q, ϵ) secure in the sense of INA-CCA2, implying

$$\mathbf{Avg\ Adv}^{\text{IND-CCA2}}_{B,\mathcal{PE}} \leq \epsilon$$

.

Thus, $\mathcal{PE}$ must be $(t, q, 2\epsilon)$ secure in the sense of NM-CCA2. □

3 Introduction to Symmetric Key Cryptography

In symmetric key cryptography, two parties share a secret key and use this key to instill certain security properties on communicated messages. Consider the case below where Alice and Bob share a secret key and wish to communicate in the presence of an adversary Eve with *unbounded computational power*.

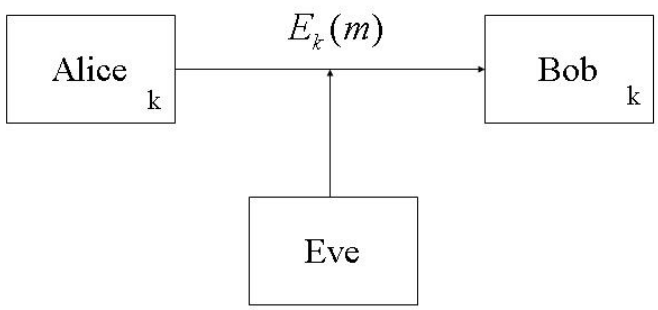

What options do Alice and Bob have to ensure security is not compromised by such a powerful adversary? Below, we will introduce the notion of *unconditional* or *perfect* security.

DEFINITION 3 *Let $E : K \times X \to Y$ denote the sample space of possible (key, message, ciphertext) triples where K, X, and Y denote the respective sample spaces on keys, messages, and ciphertexts. Let $\mathcal{K}$ and $\mathcal{X}$ be distributions over K and X, and let $\mathbf{X}$ and $\mathbf{K}$ random variables over X and K. E has prefect or unconditional security with respect to $\mathcal{X}$ if for all $y \in Y$,*

$$\Pr_{\mathcal{X},\mathcal{K}}[\mathbf{X} = x \mid E_k(x) = y] = \Pr_{\mathcal{X}}[\mathbf{X} = x]$$

DEFINITION 4 *Let $\mathcal{X}$ and $\mathcal{K}$ be independent uniform distributions over $\{0,1\}^n$. Define the one-time pad (OTP) be the following cipher: For some $k \in \mathcal{K}$, and for a single message $x \in \mathcal{X}$, let*

$$E_k(x) = k \oplus x$$

CLAIM 2
The OTP achieves perfect security.

PROOF: For all ciphertexts y,

$$\begin{aligned}\Pr_{\mathcal{X},\mathcal{K}}[\mathbf{X} = x \text{ and } E_k(x) = y] &= \Pr_{\mathcal{X},\mathcal{K}}[\mathbf{X} = x \text{ and } \mathbf{K} = x \oplus y] \\ &= \Pr_{\mathcal{X}}[\mathbf{X} = x] \Pr_{\mathcal{X},\mathcal{K}}[\mathbf{K} = x \oplus y] \\ &= \Pr_{\mathcal{X}}[\mathbf{X} = x] \frac{1}{2^n}\end{aligned}$$

Since all ciphertexts are equally likely,

$$\Pr_{\mathcal{X},\mathcal{K}}[E_k(x) = y] = \frac{1}{2^n}$$

Thus,

$$\Pr_{\mathcal{X},\mathcal{K}}[\mathbf{X} = x \mid E_k(x) = y] = \Pr_{\mathcal{X},\mathcal{K}}[\mathbf{X} = x]$$

□

CLAIM 3
If a cipher E achieves perfect security, then $|K| \geq |X|$.

PROOF: By assumption, for all $y \in Y$,

$$\Pr_{\mathcal{X},\mathcal{K}}[\mathbf{X} = x \text{ and } E_k(x) = y]$$

Recall $Y = \{E_k(x) : k \in K \text{ and } x \in X\}$. Since by construction,

$$\Pr_{\mathcal{X}}[\mathbf{X} = x] > 0 \text{ and } \Pr_{\mathcal{X},\mathcal{K}}[E_k(x) = y] > 0$$

So for all $y \in Y$,

$$\Pr_{\mathcal{X},\mathcal{K}}[\mathbf{X} = x \text{ and } E_k(x) = y] > 0$$

i.e. for all $x \in X$ and for all $y \in Y$, there exists $k \in K$ such that $E_k(x) = y$. Now suppose $|X| > |K|$. Let $y \in Y$. Then by the pigeonhole principle, there exists $x_1, x_2 \in X$, $k \in K$ such that $x1 \neq x2$ and $E_k(x_1) = E_k(x_2)$. Contradiction, because this would imply decryption is not deterministic under k. Thus, $|K| \geq |X|$. □