CS276: Cryptography September 29, 2015

CPA Security and Limitations

Instructor: Alessandro Chiesa *Scribe: Joseph Hui*

1 CPA security through PRFs

Last time we introduced the problem of security against chosen plaintext attacks. Now, using a PRF, we will construct an encryption scheme which is CPA-secure.

The encryption scheme is as follows: given a PRF $f : \{0,1\}^k \to \{0,1\}^k$, we define $E\left(1^k, sk, m\right) = (c_0, c_1) = (r, f_{sk}(r) \oplus m)$ and $D\left(1^k, sk, c\right) = f_{sk}(c_0) \oplus c_1$.

This scheme is clearly complete; we now prove CPA security, again using a hybrid argument by considering a truly random function instead of a pseudorandom one.

1 *If $\mathcal{F} = \{\mathcal{F}_k\}_k$ is a PRF, then (E, D) as described above is CPA-secure.*

Proof:

The proof proceeds in two parts.

1. First consider the uniformly random family $\mathcal{U}$. Consider $\overline{E}\left(1^k, \mathcal{U}, m\right) = r, \mathcal{U}(r) \oplus m$ and $\overline{D}\left(1^k, \mathcal{U}, c\right) = \mathcal{U}(c_0) \oplus c_1$. We claim that $\left(\overline{E}, \overline{D}\right)$ is CPA-secure. Consider any ppt distinguisher A with oracle access to $\overline{E}$ that can distinguish between two message families $\left\{m_k^{(0)}\right\}_k, \left\{m_k^{(1)}\right\}_k$. Consider $\Pr\left[A^{\overline{E}\left(1^k, \mathcal{U}_k, \cdot\right)}\left(1^k, \overline{E}\left(1^k, \mathcal{U}_k, m_k^{(b)}\right)\right) = 1\right]$ for $b = 0, 1$ (we will refer to this as $\Pr[\mathcal{A}_0]$ and $\Pr[\mathcal{A}_1]$ for convenience). Suppose that the oracle is queried at most $p(k)$ times for some polynomial $p(k)$. The i-th time the oracle is queried, it picks some random c_0 (call it $c_{i,0}$) and then returns $c_{i,0}, \mathcal{U}_k(c_{i,0}) \oplus m$. Let $\mathcal{E}$ be the event that $c_{i,0} = c_0^{(b)}$ for any i. Since each $c_{i,0}$ is chosen at random, the probability of $\mathcal{E}$ is 2^{-k} for each i and by the union bound at most $p(k)\, 2^{-k}$ overall, and this is a polynomial times an inverse exponential and therefore negligible. Now consider the case where $\mathcal{E}$ does not occur. Then what is the advantage of the algorithm, i.e. $\left|\Pr\left[\mathcal{A}_0 | \overline{\mathcal{E}}\right] - \Pr\left[\mathcal{A}_1 | \overline{\mathcal{E}}\right]\right|$? The advantage is necessarily zero: because $c_{i,0} \neq c_0^{(b)}$ for all i, and $\mathcal{U}$ is uniformly random, the value of b is independent from all of the data that A has seen. That is to say, for every scenario where A answers correctly, the scenario where A answers incorrectly, on the exact same inputs and answers from the oracle, occurs with the same probability, and A answers correctly with probability precisely $\frac{1}{2}$.
   So, following the total probability law, we have

$$\left|\Pr[\mathcal{A}_0] - \Pr[\mathcal{A}_1]\right| = \left|\Pr[\mathcal{E}]\left(\Pr[\mathcal{A}_0|\mathcal{E}] - \Pr[\mathcal{A}_1|\mathcal{E}]\right) + \Pr\left[\overline{\mathcal{E}}\right]\left(\Pr\left[\mathcal{A}_0|\overline{\mathcal{E}}\right] - \Pr\left[\mathcal{A}_1|\overline{\mathcal{E}}\right]\right)\right|$$

   which is at most $\left|\text{negl} \cdot 1 + \Pr\left[\overline{\mathcal{E}}\right] \cdot 0\right|$, and therefore negligible.

2. Now we will show that the pseudorandom function "looks like" a truly random function. Let $\{m_k\}_k$ be any family of messages and consider

$$\left|\Pr\left[A^{\overline{E}\left(1^k, \mathcal{U}_k, \cdot\right)}\left(1^k, \overline{E}\left(1^k, \mathcal{U}_k, m_k\right)\right) = 1\right] - \Pr\left[A^{E\left(1^k, sk, \cdot\right)}\left(1^k, E\left(1^k, sk, m_k\right)\right) = 1\right]\right|$$

(which we will refer to as $|\Pr[\mathcal{A}] - \Pr[\mathcal{A}']|$ for convenience). Suppose towards a contradiction that there exists some ppt distinguisher A with oracle access which distinguishes between these two distributions with non-negligible advantage. We will construct a ppt distinguisher $B^{\mathcal{O}}$ which distinguishes between $\mathcal{F}$ and $\mathcal{U}$. $B^{\mathcal{O}}$ works as follows:
First, sample $r \leftarrow \{0,1\}^k$.
Second, construct a challenge for A: $c \leftarrow (r, \mathcal{O}(r) \oplus m)$.
Third, simulate $A\left(1^k, c\right)$, answering each query m_i with $(r_i, \mathcal{O}(r_i) \oplus m_i)$ for i.i.d. r_i.
Finally, output whatever A says.
Now, whenever B is given $\mathcal{O} = \mathcal{F}$, then A is exactly given $E\left(1^k, sk, \cdot\right)$, and vice versa. So, the probability that B succeeds is exactly the probability that A succeeds. So if A has non-negligible advantage, so does B, a contradiction to the pseudorandomness of $\mathcal{F}$. Thus the probability $|\Pr[\mathcal{A}] - \Pr[\mathcal{A}']|$ must be negligible.

Now, we want to show that $|\Pr[\mathcal{A}'_0] - \Pr[\mathcal{A}'_1]|$ is negligible, which is the definition of CPA security. This is equal to

$$|\Pr[\mathcal{A}'_0] - \Pr[\mathcal{A}_0] + \Pr[\mathcal{A}_0] - \Pr[\mathcal{A}_1] + \Pr[\mathcal{A}_1] - \Pr[\mathcal{A}'_1]|$$

, and by the triangle inequality, this is at most

$$|\Pr[\mathcal{A}'_0] - \Pr[\mathcal{A}_0]| + |\Pr[\mathcal{A}_0] - \Pr[\mathcal{A}_1]| + |\Pr[\mathcal{A}_1] - \Pr[\mathcal{A}'_1]|$$

, all of which are negligible (by the previous two steps). So this probability is indeed negligible and we are done. □

2 Limitations

We will discuss some limitations of this scheme. Firstly, we have ciphertext expansion: the ciphertext is larger than the message. Secondly, we operate over a fixed message space: we cannot encode arbitrarily long messages this scheme in the naive way. Finally, this is CPA-secure, which is quite strong, but there still remain more elaborate attacks that CPA security does not defend against.

The second limitation is easy to deal with: if we have a vector of messages to send instead of just one, we just generate a fresh random r for each message.

2.1 Modes of operation

We can also deal with the first limitation through an idea called "modes of operation". Given a PRF/PRP f, we will construct a scheme that is CPA-secure, with arbitrary message lengths and expansion factor ≈ 1. We'll consider some historical attempts at this.

One mode of operation is electronic code book mode (ECB), where for a PRP f, we have

$$E\left(1^k, sk, \overrightarrow{m}\right) = f_{sk}(m_1), f_{sk}(m_2), ..., f_{sk}(m_l)$$

. This has the advantage that it is parallelizable (each component can be computed separately). However, this is obviously insecure against CPA, because it is deterministic: the adversary can check any decryption just by asking the decryption oracle.

Another is cipher block chaining mode (CBC). Given a PRP f, we pick r and generate a ciphertext $c_1 = f_{sk}(r) \oplus m_1$. Then with the next component of the message, we generate the next ciphertext

$c_2 = f_{sk}(c_1) \oplus m_2$, and so on. Finally we output $(r, c_1, c_2, ..., c_l)$. This is CPA-secure if f is a PRP. Unfortunately, the encryption cannot be computed in parallel as before. On the other hand, decryption can be computed in parallel (and with random access, i.e. one can choose to just decrypt the 52nd component of the message). The expansion ratio is $1 + \frac{1}{l}$.

A third is output feedback mode (OFB), which uses a PRF f. We won't discuss it here; we will just mention that it is sequential but encryption can be preprocessed (similarly to the next mode).

Finally, we have randomized counter mode (CTR). Here, we generate a random r as in the original scheme, and the encryption function is

$$E\left(1^k, sk, \overrightarrow{m}\right) = f_{sk}(r+1) \oplus m_1, f_{sk}(r+2) \oplus m_2, ..., f_{sk}(r+l) \oplus m_l$$

Not only are encryption and decryption parallelizable, but given the key, one can compute the "one-time pad" in advance and just XOR the message when it is available, making encryption incredibly efficient. The expansion ratio is again $1 + \frac{1}{l}$, and it is CPA-secure given a PRF.

2.2 Malleability attacks

Now we will consider some attacks that CPA security alone does not prevent. In this scenario, we have an active attacker who does not only read messages but also modifies them. Now, let's consider the CPA-secure encryption scheme we described. Because the encryption is simply the XOR-ing of the message with some pseudorandom bits, the attacker (Eve) can modify the message by flipping some of the bits. Indeed, if she knows some of the bits, she can then set those bits to whatever values she likes. For example, suppose she has asked Alice to send a 32-bit message to Bob, which signifies a dollar amount to be paid to Eve. If Eve knows that Alice will indicate a relatively meager sum (perhaps a few hundred dollars), she can flip the first bit of the ciphertext before passing it on to Bob, thereby increasing her profits by $2,147,483,648. We would like to avoid these kinds of attacks in which an attacker can turn one well-formed ciphertext into another.

Additionally, suppose that Eve can selectively discover decryptions of certain messages. Certainly we cannot allow her to decrypt arbitrary messages, or no encryption scheme could possibly succeed. However, it's possible that Eve can (say) decrypt some messages but not others. We might imagine, for example, that foolish Bob has carelessly left his computer unlocked while he is out to lunch, and clever Eve is able to test decryptions of various ciphertexts for a short period of time, in preparation for trying to decrypt some unknown message that will be sent later on. In the most extreme case, we could suppose that Eve can in fact obtain the decryption of every message except the message that Alice sends. In this case, our CPA-secure scheme also fails: if Eve is challenged on (c_0, c_1), she can ask for the decryption of $(c_0, 0)$, thereby obtaining $f_{sk}(c_0)$, which she can XOR with the ciphertext to recover the message. We will consider how to build a scheme that can resist even this attack (called a chosen-ciphertext attack).

References