

Lecture 7 - CCA security

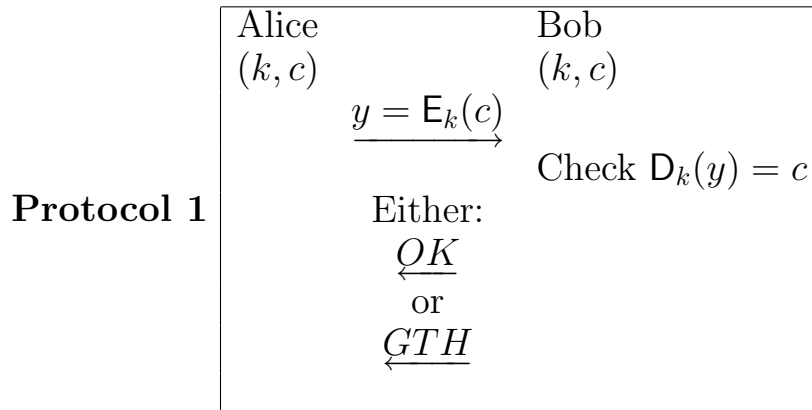
Boaz Barak

October 9, 2007

Reading KL Section 3.7 (pages 103–104), Chapter 4 (although we follow a somewhat different order).

The online shopping problem Let's abstract the online shopping problem as follows: Alice and Bob share a key k , and Alice also has a credit card number c which Bob can validate against a database. (For simplicity, assume that there is only one valid credit card, and so the validation is simply checking that the card equals c .)

Consider the following protocol for an online transaction:



Security What does it mean for this protocol to be secure. A security definition always has three components:

- The capabilities of the adversary: we'll always model it as a polynomial-time algorithm.
- The attack model: in this case we'll start with a passive attack - Eve just sits on the line and listens to polynomially many executions of Protocol 1 (all using the same key and credit card).
- The definition of a break / "win" of the adversary: We'll make the following definition: Assume that $k \leftarrow_{\text{R}} \{0, 1\}^n$ and $c \leftarrow_{\text{R}} [10^8]$. We say

that the protocol is secure if for every poly-time Eve, poly-bounded ϵ and large enough n ,

$$\Pr[\text{Eve guesses } c] \leq 10^{-8} + \epsilon(n)$$

We say that a protocol satisfying the above definition is a *passively secure authentication protocol*

Theorem 1 If (E, D) is CPA-secure then Protocol 1 is a passively secure authentication protocol.

Proof We'll do the proof by reduction: assume that A can break the passive security of Protocol 1, we'll build an adversary B that can break the CPA security of the encryption scheme.

In an actual attack, A sees T repetitions of the following messages:

$$\begin{array}{c} \xrightarrow{E_k(c)} \\ \xleftarrow{OK} \end{array}$$

(of course each occurrence of $E_k(c)$ uses fresh randomness)

Consider the following “mental experiment” we run the attack of A , but instead choose a junk message $J \leftarrow_{\text{r}} \{0, 1\}^n$ (e.g., $J = 0^n$) and send an encryption of J instead of c . That is, A will see T repetitions of:

$$\begin{array}{c} \xrightarrow{E_k(J)} \\ \xleftarrow{OK} \end{array}$$

Clearly in this case A gets no information on c and so will guess it successfully with probability at most 10^{-8} . Hence under our assumption A can distinguish between the two cases with advantage ϵ . Using the hybrid method it follows that there is an $i \in [T]$ such that A can distinguish between the case that the first i messages are encryptions of c and the last $T - i$ messages are encryption of J and the case that the first $i + 1$ messages are encryptions of c and the last $T - i - 1$ messages are encryptions of J . Specifically, if we let p be the probability that A guesses correctly in the first case then in the second case A guesses correctly with probability at least $p + \epsilon/T$.

We can now mount a CPA attack against the encryption scheme. B will do the following:

1. Choose a number $c \leftarrow_{\text{r}} [10^8]$, $J \leftarrow_{\text{r}} \{0, 1\}^n$.
2. Ask for i encryptions of c .

3. Choose $x_0 = J, x_1 = c$ and give x_0, x_1 to the sender to obtain $y = E_k(x_i)$ for $i \leftarrow_R \{0, 1\}$.
4. Ask for $T - i - 1$ encryptions of J .
5. Feed all the encryptions plus “OK” messages to A and get a guess g from A .
6. If $g = c$ then output 1; otherwise output a random bit in $\{0, 1\}$.

We have that if $i = 0$ then we output 1 with probability $p + (1 - p)\frac{1}{2} = \frac{1}{2} + \frac{p}{2}$, while if $i = 1$ we output 1 with probability at least $\frac{1}{2} + \frac{p}{2} + \frac{\epsilon}{2T}$. Hence, the probability we make the correct guess is at least

$$\frac{1}{2}\left(\frac{1}{2} + \frac{p}{2} + \frac{\epsilon}{2T}\right) + \frac{1}{2}\left(1 - \frac{1}{2} - \frac{p}{2}\right) = \frac{1}{2} + \frac{\epsilon}{4T}$$

□

Active attacks In an active attack we allow Eve to the drop, insert and modify messages between Alice and Bob. (Of course, she can completely cut off the communication between them, but then she won’t learn the credit card number.) Alice and Bob can always check that messages are of the proper length etc., and so we’ll assume that Eve only sends well formed messages (though this is a huge issue in practice!).

Does the proof go through? To simulate Eve in this case, we’d like to change Step 5 above to the following: For every generated encryption y_i , feed y_i to A and obtain the modified version y'_i . Then, check if y'_i is an encryption of c and if so, feed “OK” to A ; otherwise feed it with GTH. But of course, we cannot check if y'_i is an encryption of c in the course of running an attack on the encryption scheme!

Is it just the proof? It seems “obvious” that the problem is with the proof method and that Protocol 1 is secure even against active attacks - after all how can Eve guess the credit card when it’s never sent in the clear??

Counterexample It turns out the intuition is false - there is a CPA secure encryption which makes Protocol 1 *insecure* against active attacks! The counterexample below may seem contrived (and it is) but it illustrates an idea that was used in actual attacks on real-life protocols (specifically SSL V3.0).

Consider the unary encoding of decimal digits as bit strings: the digit i is encoded a string of 11 bits starting with i zeros and then a 1. To be more specific, the encoding algorithm will encode i as $0^i 1^{11-i}$ but the

decoding algorithm will only check for existence of a single 1 following the first i zeroes.

Now suppose that the credit card number c is encoded in this form to a string \hat{c} , and then we use the PRF-based CPA-secure encryption we saw in class ($E_k(x) = (r, f_k(r) \oplus x)$) to encrypt \hat{c} .

The crucial observation is that if Eve gets a ciphertext $E_k(x) = (r, y)$ even without knowing what is the plaintext, she can still flip the i^{th} bit of y and that will change the ciphertext to the encryption of x with its i^{th} bit flipped.

This means that Eve to guess a particular digit of the card, Eve can try to flip the last location in the encoding of that digit, then one before last, and so on. If i is the first location where Eve got a “GTH” message then the digit was $i - 1$.

Thus, after at most 88 executions Eve would be able to learn the entire credit card number.

Stronger notion of encryption: It turns out there is a stronger notion of encryption scheme, that suffices to imply security of Protocol 1 and in many other places. This notion might seem “crazy strong” but it turns out that:

- It is needed for many applications.
- There are pretty efficient constructions conjectured to satisfy it.

Chosen ciphertext attack The definition of chosen ciphertext attack is a variation of chosen plaintext attack but this time the attacker is given oracle access even to the decryption box!

- Adversary gets oracle access to $x \mapsto E_k(x)$ and $y \mapsto D_k(y)$.
- Adversary chooses x_0, x_1 .
- Sender chooses $k \leftarrow_R \{0, 1\}^n$, $i \leftarrow_R \{0, 1\}$ and sends $y^* = E_k(x_i)$ to the adversary.
- Adversary gets access to $x \mapsto E_k(x)$ and $y \mapsto D_k(y)$ with one restriction - it cannot ask the decryption box the exact string y^* .
- Adversary comes up with a guess j . She *wins* if $i = j$.

(E, D) is CCA-secure if for every poly Adv , poly-bounded $\epsilon : \mathbb{N} \rightarrow [0, 1]$, and large enough n , $\Pr[\text{Adv wins}] \leq 1/2 + \epsilon(n)$.

Theorem 2 If (E, D) is CCA-secure then Protocol 1 is an actively secure authentication protocol.

Proof We modify the proof of Theorem 1 in the following way: first, we assume that Bob sends an OK message if the encryption is either of c or the junk message J . We can do that because the probability that Eve sends an encryption of the junk message J is at most $\frac{T}{2^n}$.

Then, we can simulate Eve - in Step 5 we do the following: for every generated encryption y_i , feed y_i to A and obtain the modified version y'_i . Then, if $y'_i = y^*$ then feed "OK" to A ; otherwise, ask the decryption box to decrypt y'_i and answer OK iff it decrypts to either c or J .

Construction of CCA secure encryption scheme The following encryption scheme is CCA secure: let $\{p_k\}$ be a collection of pseudorandom permutation mapping $\{0, 1\}^{3n}$ to $\{0, 1\}^{3n}$:

- To encrypt $x \in \{0, 1\}^n$ do the following: choose $r \leftarrow_{\text{R}} \{0, 1\}^n$, and send $p_k(x||r||0^n)$ (where $||$ denotes concatenation).
- To decrypt $y \in \{0, 1\}^{3n}$, compute $x||r||w = p_k^{-1}(y)$. if $w \neq 0^n$ then output \perp . Otherwise, output x .

Something to think about for Thursday Is there a CPA secure encryption where an active Eve can set every bit in the encrypted text to an arbitrary value of her choice + flip every bit she wants to?