Lecture 21 - Forward Security, Identity-Based Encryption

Boaz Barak

December 6, 2005

Forward security. We said that protecting secret keys is crucial for cryptography, and gave some techniques to protect the key in the previous class.

But what happens if the adversary does learn the key? Indeed, suppose I have a secret decryption key, corresponding to some known public key which I use for a long time, and at some point an attacker breaks into my computer and learns the secret key without my learning all about it.

It's clear that from now on, the attacker will be able to read all the encrypted messages that are sent to me. It's also seems intuitively clear that if the attacker recorded previously the ciphertexts of the encrypted messages that were sent to me before he gained access to my computer, then now he will be able to use my secret key to decrypt these messages. Surprisingly (or perhaps not, since this is crypto and strange things always happen) this intuition is false, and it is possible to ensure that the attacker will only be able to decrypt message sent after he broke into my computer, even if I don't know when or whether or not he broke into it.

Forward-secure encryption schemes Encryption schemes that maintain this property are called *forward secure*. A forward secure public key encryption scheme has the following components:

Key generation As usual, G outputs a public key e and a secret key, which we denote by d_0.

Encryption algorithm The encryption algorithm E takes as usual as inputs the encryption key e and the message to be encrypted m. However, it takes also an additional input t, which is the current time (or time period).

Update mechanism We'll use a different secret key to decrypt at each time period t. We'll denote this secret key by d_t. Of course, we must be able to efficiently compute d_t for every t given the original key d_0 (since that is all the information the receiver has). However, we'll actually do it in the following way: there is an algorithm $UPDATE$ that on input t and d_{t-1} outputs d_t. At the beginning of each time period t, the receiver will compute $d_t = UPDATE(t, d_{t-1})$ and *erase* d_{t-1} from its memory.[1] (It will be the case that this algorithm $UPDATE$ is hard to invert, that is, from d_t it's hard to come up with d_{t-1}.)

Decryption To decrypt a message sent at time t, we'll use d_t. Thus the validity condition is that for every m, $\mathsf{D}_{d_t}(\mathsf{E}_e(m,t)) = m$.

[1] Note that there are many technical difficulties involved in securely erasing memory from modern computers, that use hard-drives, virtual memory and paging. We ignore these issues here.

We can define forward-secure variants of both CPA security and CCA security. The idea is that we run the usual attack game (either CPA or CCA), except that there is a global time counter t that the adversary can ask to increase by one from time to time. The adversary then chooses two messages m_1 and m_2 and gets the challenge — an encryption of m_b for $b \leftarrow_R \{1,2\}$. The game again continues as in the usual CPA/CCA case. we continue this game as usual. However, before the adversary needs to guess b, the time counter t is increased by one, and the adversary is given the secret key d_t. The adversary can then use that information in order to try to guess b with probability greater than $1/2$.

Other forward secure primitives The notion of forward security is pretty general, and there are definitions and constructions for forward secure signature schemes, pseudorandom generators, private key encryption, and others.

Constructing forward secure encryption schemes We are going to construct forward secure encryption schemes using another object that is called *identity-based encryption schemes*. Identity-based encryption schemes are themselves just as fascinating (and perhaps even more) as forward-secure encryption schemes. The idea was first suggested by Shamir in the 80's but a construction was only given in 2001 by Boneh and Franklin. Even that construction was only proven secure in the random oracle model and getting a random-oracle free construction seemed to be a very hard problem to many researchers (including myself). However in 2004, Boneh and Boyen (improving on Canetti, Halevi and Katz) managed to get a construction proven secure under reasonable computational assumptions, without any random oracles.

Identity based encryption The idea of identity based encryption is very simple - what if your name could be your public key? That is, where in standard public-key crypto, if I want to send Dave a secure email he has to send me his public key (or perhaps publish it in a public key directory) in IBE my encryption algorithm simply takes the string "Dave Xiao" as an input.

More accurately, an identity-based encryption (IBE) is comprised of the following parts:

Master key generation There is an algorithm G_{master} that generates the master public and private keys pub_{master} and $priv_{master}$.

Key derivation There is an algorithm $Derive$ that gets as input the private key $priv_{master}$ and an arbitrary string $id \in \{0,1\}^*$, and outputs a decryption key d_{id}.

Encryption To encrypt a message m to ID id, run $\mathsf{E}_{pub_{master},id}(m)$.

Decryption There is a decryption algorithm that takes as input the decryption key k_{id} and a ciphertext y, where the validity condition is that for every m and id,

$$\mathsf{D}_{d_{id}}(\mathsf{E}_{pub_{master},id}(m)) = m$$

Again IBE can be defined with either a CPA or CCA variant. In both cases the adversary gets the public master keys pub_{master} and runs the usual CCA/CPA attack. However, it now gets an additional oracle access to the key derivation algorithm, to which it can query a string id and get back d_{id}. When making the challenge the adversary not only specifies two messages m_1 and m_2 but also an ID id^*, and gets the challenge ciphertext $y^* = \mathsf{E}_{pub_{master},id^*}(m_b)$ for $b \leftarrow_R \{1,2\}$. The adversary has now additional access to key derivation algorithm, but conditioned on not asking the query id^*, and if it's a CCA attack also access to the decryption oracle, where it can make any query of the form $\langle id, y\rangle$ as long as either $id \neq id^*$ or $y \neq y^*$. Again, the adversary is successful if it guesses b with probability noticably higher than $1/2$.

Forward-secure encryption from IBE Given an IBE scheme, one can construct a forward-secure encryption in the following way:

- The public and private keys are generated as follows: generate pub_{master} and $priv_{master}$ using the generator for the IBE scheme. Assuming we're going to use this scheme for T time period, for every $1 \leq t \leq T$, let id_t denote the string "time slot t" and let k_t denote $DERIVE(priv_{master}, id_t)$. The private key d_0 will be the concatenation of k_1 until k_T.
- To encrypt at time t simply run $\mathsf{E}_{pub_{master},id_t}(m)$.
- The key d_t will be the concatenation of k_t until k_T. That is, the update mechanism at time t involves erasing the key k_t from the list.

It's not hard to prove this scheme is secure. However, its drawback is that it requires the private key to be of size nT and maintaining such a large secret storage may be infeasible. It can be easily improved however, to require the receiver to only public storage of this length: instead of storing $k_1, \ldots, k_T$ store k_1 in private and keep a non secret file $y_2, \ldots, y_T$ where $y_t = \mathsf{E}_{pub_{master},k_{t-1}}(y_{t-1})$. There are also constructions without need for any storage that depends on T worse than logarithmically.

Other applications IBE has several other potential applications. For example, suppose that when sending email to me, people use the ID "Boaz Barak $\circ$ current date". Then, when I go to a conference with my laptop, I can keep in the laptop only the private keys corresponding to these dates. It has also been suggested that a manager can use IBE to provide assistants with "restricted private keys" that can only decrypt messages with particular subjects. In any case IBE is quite cool. In fact, Boneh and others formed a company (Voltage) based on the IBE technology.

Construction of IBE We present the random oracle based construction of Boneh and Franklin.

Pairing diffie hellman assumption The DDH assumption says that in an appropriate cyclic group G with a generator g, it is impossible to distinguish between the triple (g^a, g^b, g^{ab}) and the triple (g^a, g^b, g^c) for x, y, z chosen independently at random from $\{0, \ldots, |G| - 1\}$.

Consider the following question: can there be a group where it's actually *easy* to compute g^{ab} from g^a, g^b but given g^a, g^b, g^c it's *hard* to compute g^{abc}. It's not hard to see that this is impossible - if you have an algorithm to compute the first problem, you can apply it twice to obtain first g^{ab} and then g^{abc} to solve the second problem. However, we will somehow manufacture a situation where this is "morally true". We are going to consider two cyclic groups G and H with $|G| = |H|$ and generators g and h respectively and a function $f : G \times G \to H$ satisfying the following: $f(g^a, g^b) = h^{ab}$. It turns out it is possible to come up with such groups and a function. In some sense we manage to solve the first problem, but only when moving to a different group.

We are going to make the following assumption (called **pairing DDH**): for random a, b, c, d it is impossible to distinguish between $\langle g^a, g^b, g^c, f(g^a, g^b)^c = h^{abc} \rangle$ and $\langle g^a, g^b, g^c, h^d \rangle$.

Assuming this, we will build an identity-based cryptosystem as follows:

Public and private master keys Generate groups G, H and generators g, h and function $f : G \times G \to H$ such that $f(g^a, g^b) = h^{ab}$. Let $R : \{0,1\}^* \to G$ denote a random oracle. Choose a at random from $\{0, \ldots, |G| - 1\}$ and publish g^a. a is the secret key.

Identity keys For an identity id, let $e_{id} = R(id)$ (the random oracle applied to the string id. We let $b \in \{0, \ldots, |G| - 1\}$ be a number such that $e_{id} = g^b$. Note that no one (including even the holder of the master private key) knows b. The secret key for id, $d_{id} = e_{id}^a = g^{ab}$. Note that it can be derived using the private key.

Encryption To encrypt a message m for ID id, choose $c \leftarrow_{\mathrm{R}} \{0, \ldots, |G| - 1\}$, compute $\pi = f(g^a, e_{id})^c = h^{abc}$ and send $g^c, \pi \oplus m$.

Decrypt Given the secret key g^{ab} and the message $g^c, h^{abc} \oplus m$, the receiver computes $\pi = f(g^{ab}, g^c) = h^{abc}$ and uses that to retrieve the message.

Assuming that g^d for a random d is represented as a random string, this scheme can be shown to be CPA secure under the pairing DDA assumption. By further using (or abusing?) the random-oracle it can be shown secure under a weaker assumption (namely pairing computational Diffie-Hellman (CDH) assumption) and can also be made CCA secure. For more details of the proof, see the paper by Boneh and Franklin. **Note:** In that paper, as in most other papers in this subject, *additive* notation is used for the group G (but not for H). Thus, instead of g^a you will see there $a \cdot g$, and the pairing DDH assumption will be that for random a, b, c, d the tuple $\langle a \cdot g, b \cdot g, c \cdot g, f(a \cdot g, b \cdot g)^c \rangle$ is indistinguishable from $\langle a \cdot g, b \cdot g, c \cdot g, f(a \cdot g, b \cdot g)^d \rangle$