Advanced Topics in Cryptography

Lecture 11: Chosen-ciphertext security from identity based encryption.

Benny Pinkas

An announcement

- Seminar talk, next Wednesday:

  Hovav Shacham

  New paradigms in signature schemes
- Abstract:
  - Groups featuring a computable bilinear map are particularly well suited for signature-related primitives.
  - For some signature variants the only construction known is based on bilinear maps.
  - Bilinear-map-based constructions are simpler, more efficient, and yield shorter signatures.
  - The talk describes three constructions and their applications: short signatures, aggregate signatures, group signatures.

Related papers

– Chosen-Ciphertext Security from Identity-Based Encryption. D. Boneh, R. Canetti, S. Halevi, and J. Katz.

– http://crypto.stanford.edu/~dabo/papers/ccaibejour.pdf

Chosen-ciphertext security

- Chosen-plaintext security (CPA)
  - Semantic security
  - Indistinguishability
- CPA does not protect against active attacks
- Chosen-ciphertext security (CCA)
  - The adversary can get decryptions of ciphertexts of his choice
  - This is the *de facto* required level of security today.
  - *Non-adaptive CCA:* adversary can ask decryption queries before receiving its challenge
  - *Adaptive CCA:* adversary can ask decryption queries even after receiving its challenge

Security against chosen-ciphertext attacks

- The game:
  - We show the public key to the adversary
  - Adversary can ask to receive decryptions of messages of his choice
  - Adversary chooses two messages m_0,m_1 (possibly based on the answers he previously received)
  - Adversary is given an encryption $E(m_b)$, where $b \in_R \{0,1\}$
  - Adversary can issue further decryption queries, but not $E(m_b)$ *(this is the difference between adaptive and non-adaptive attacks)*
  - Adversary guesses b

- Adversary succeeds if its probability of guessing b correctly is not negligibly close to ½

CCA-secure encryption schemes

- Constructions based on the random oracle model (OAEP and its variants)
- Generic constructions
  - Based on a CPA-secure encryption scheme and non-interactive zero-knowledge proofs (NIZK).
  - Show feasibility.
  - Not very practical. NIZK proofs are based on reductions to NP-complete problems.
- Algebraic constructions
  - Cramer-Shoup.
  - Based on the DDH and similar problems.

New construction

- A CCA-secure public encryption scheme
  - Based on a generic assumption: the existence of a CPA-secure identity based encryption scheme.
  - Specific instantiations, based on number theoretic assumptions, can be almost as practical as Cramer-Shoup.
  - Unlike previous CCA-secure schemes, does not use a "proof of well formedness".

Identity based encryption (IBE)

- A public-key encryption scheme where the key can be an arbitrary string
- Key generation center (KGC)
  - Holds the master private key
  - Generates public system parameters
- Key derivation: The KGC can provide each user with the private key corresponding to his/her name.
  - The private key is a function of the name (or an arbitrary string) and the master private key
- Encryption: everyone can encrypt messages to Alice. The ciphertext is a function of the plaintext, Alice's name, and the public parameters.
- Decryption: Alice uses her private key and the system parameters to decrypt messages sent to her

IBE – security definitions

- Main challenge: adversary can get private keys of some identities, while attacking a different identity
- Adaptively-chosen-key semantic (CPA) security

1. The adversary obtains keys for a polynomial number of IDs, which it chooses adaptively
2. It outputs a different ID*, and two messages m_0,m_1
3. It receives E(m_b,ID*), for $b \in_R \{0,1\}$
4. The adversary tries to guess b

- Selective-ID IBE
  - A weaker notion of IBE
  - The adversary must select ID* before receiving the IDs in Step 1 (i.e., ID* is not a function of Step 1).

Identity based encryption

- Master Key Generation:
  - $MKG(1^k) \rightarrow (PK_{master}, SK_{master})$
- Key Generation:
  - $G(ID, SK_{master}) \rightarrow SK_{ID}$
- Encryption:
  - $E(m, ID, PK_{master}) \rightarrow c$
- Decryption
  - $D(c, ID, SK_{ID}) \rightarrow m$ such that $c = E(m, ID, PK_{master})$

The construction

- Based on
  - An IBE scheme with chosen-plaintext selective-ID security (even weaker than full pledged IBE)
  - A one-time signature scheme
    - Each key is used only for a single signature
    - Strong unforgeability: the adversary should not forge a new signature even on a previously signed message

- Key generation:
  - The user runs the master key generation algorithm of the IBE scheme, $MKG(1^k) \rightarrow (PK_{master}, SK_{master})$. Its public key is PK_{master}.

The construction

- Encryption: to encrypt m,
  - The sender generates fresh signing and verification keys for the signature scheme, *sk, vk*.
  - The sender encrypts *m* with respect to the identify *vk*. $E(m,vk,PK_{master}) \rightarrow c$
  - It signs the resulting IBE ciphertext $sign_{sk}(c) \rightarrow \sigma$.
  - The ciphertext is $\langle vk,c,\sigma \rangle$.
- Decryption of $\langle vk,c,\sigma \rangle$:
  - The receiver uses *vk* to verify that σ is a signature of c. If not, it aborts.
  - The receiver computes the IBE private key $G(vk,SK_{master}) \rightarrow SK_{vk}$.
  - It then computes the decryption $D(c,vk,SK_{vk}) \rightarrow m$.

Security:

- Warmup: security against *non-adaptive* CCA attacks
  - Instead of using signatures, the sender
    - Chooses a random string *r*
    - Uses the IBE scheme to encrypt *m* under the identity *r,* resulting in a ciphertext *c*.
    - Sends $\langle r,c \rangle$ to the receiver.
  - The receiver decrypts *c* using the secret key of ID *r*.
- Security of this variant:
  - The adversary can only do decryption queries *before* receiving the challenge ciphertext. That is, before learning the value *r* of the ciphertext it has to break.
  - Therefore, it uses different *r* values in its queries.
  - The IBE scheme is secure even if the adversary learns the decryption keys of many IDs *r'*, different than *r*.

Security - intuition

- Say that a ciphertext $\langle vk,c,\sigma \rangle$ is valid if the verification key *vk* verifies that σ is a signature of c.
- The adversary is given a challenge ciphertext $\langle vk^*,c^*,\sigma^* \rangle$
- Suppose that the adversary submits a ciphertext $\langle vk,c,\sigma \rangle \neq \langle vk^*,c^*,\sigma^* \rangle$ for decryption
  - If vk=vk*, then $\langle vk,c,\sigma \rangle$ cannot be valid (this would have meant that the adversary generated a new signature pair (c,σ), even though it does not the signature key).
  - Therefore vk$\neq$vk*. The selective-ID security of the IBE scheme implies that a decryption of c (and even the decryption key for the identity *vk*), do not compromise encryptions done with the id *vk**.

Security proof

- THM: if the IBE scheme is selective-ID secure against chosen-plaintext attacks, and the signature has strong one-time security, then the system has CCA security against adaptive attacks.
- Proof:
  - Assume that A attacks the system in an adaptive CCA attack, and is given the challenge ciphertext $\langle vk^*,c^*,\sigma^*\rangle$.
  - Let FORGE denote the event that A submits a valid ciphertext $\langle vk^*,c,\sigma\rangle$ to the decryption oracle ($c,\sigma \neq c^*,\sigma^*$).
  - Claim 1: The probability of FORGE is negligible.
  - Claim 2: $|\ Pr(Success\ \&\ \neg FORGE) + 0.5Pr(FORGE) - 0.5|$ is negligible.

Why this proves the theorem

- $|\Pr(\text{Success}) – 0.5)|$

- $\leq |\Pr(\text{Success} \;\&\; \text{FORGE}) – 0.5\Pr(\text{FORGE})| +$
  $|\Pr(\text{Success} \;\&\; \neg\text{FORGE}) + 0.5\Pr(\text{FORGE}) – 0.5|$

- $\leq \Pr(\text{FORGE}) +$
  $|\Pr(\text{Success} \;\&\; \neg\text{FORGE}) + 0.5\Pr(\text{FORGE}) – 0.5|$

Proof of Claim 1

- The probability of FORGE is negligible
- Proof:
  - We construct a forgery algorithm F for the signature scheme which can forge signatures with probability Pr(FORGE).
  - F has access to a signature algorithm, which is willing to sign a single message.
  - F is given a verification key *vk**. It generates the public key of the IBE system, and provides it to the adversary A.
  - F can answer any decryption query of A.
  - When A provides F with m_0, m_1, F chooses $b \in_R \{0,1\}$, encrypts m_b with the ID vk*, and asks for a signature $\sigma*$ on this ciphertext c*. It returns $\langle vk^*, c^*, \sigma* \rangle$ as the challenge.
  - If A submits a ciphertext $\langle vk^*, c, \sigma \rangle$, F obtained a forgery.

Proof of Claim 2:

$|\Pr(\text{Success} \,\&\, \neg\text{FORGE}) + 0.5\Pr(\text{FORGE}) - 0.5|$ is negligible

- We construct A' which attacks the IBE scheme:
  - A' generates *(vk*,sk*)* and sets the target ID to *vk**. A' is given a master public key *PK* (to attack) and sends it to A.
  - A makes a decryption query $\langle vk,c,\sigma\rangle$.
    - If *vk*=*vk**, and the signature σ is good, A' aborts.
    - If the signature σ is incorrect, A' returns "fail".
    - If $vk \neq vk^*$, and the signature σ is good, A' asks for SK_{vk}, and uses it to decrypt c and return the answer to A.
  - A sends m_0,m_1 to A'. A' sends them to its decryption oracle, with the ID *vk**. It receives an encryption c* of m_b, signs it and sends the answer $\langle vk^*,c^*,\sigma^*\rangle$ to A.
  - A' continues as before. When A outputs b', A' outputs b=b'.
- A' is a perfect simulation for A, except in case of forgery:
  - $|\Pr_{A'}(\text{Success})-0.5| = |\Pr_A(\text{Success} \,\&\, \neg\text{FORGE})+0.5\Pr_A(\text{FORGE})-0.5|$

One time signatures

- Signature scheme for a single message

- Example: to sign a single bit
  - Private signature key: $x_0, x_1 \in \{0,1\}^k$
  - Public verification key: $h_0=h(x_0), h_1=h(x_1)$, where *h* is one-way
  - Signature (of bit *b*): x_b
  - Verification: check that $h(x_b) = h_b$

- Very efficient
- Given signature of *b*, adversary cannot fake a signature of *1-b*

One time signatures

- Signing message of size *n:*
  - Private key: $\{ x_{i,0}, x_{i,1} \}_{i=1..n}$
  - Public key: $\{ h(x_{i,0}), h(x_{i,1}) \}_{i=1..n}$
  - Signature of $b_1,\ldots,b_n$: $x_{1,b1},\ldots,x_{n,bn}$
- Alternatively,
  - Private key: $\{ x_i \}_{i=1..n+log(n)}$
  - Public key: $\{ h(x_i) \}_{i=1..n+log(n)}$
  - Signature of $b_1,\ldots,b_n$: x_j for all $b_j=0$. Let $c_1,\ldots,c_{\log(n)}$ be the Hamming weight of b. Open also x_{n+j} for all $c_j=0$.
  - Very efficient
    - Can use a full signature scheme to sign public key of one-time scheme (offline).
    - When it is required to sign *m*, signing can be done very efficiently.
  - What happens if two different messages are signed with the same public key?

A construction of selective-ID IBE with no random oracle assumptions

One-time signatures