Advanced Cryptographic Primitives: Lecture 7

Scribe: François Pirot

M2IF

0.1 Applications of (H)IBE to chosen-ciphertext security

0.1.1 Definition

Definition (Rackoff-Simon, Crypto'91 [1]) A public-key encryption scheme is secure against adaptive chosen ciphertext attacks (IND-CCA2) if no PPT adversary $\mathcal{A}$ has non-negligible advantage on the following game:

1. The challenger generates $(PK, SK) \leftarrow \text{Keygen}(\lambda)$ and gives PK to $\mathcal{A}$
2. $\mathcal{A}$ invokes the decryption oracle a polynomial number of times: at each query, $\mathcal{A}$ chooses a ciphertext C and obtains $M \leftarrow \text{Decrypt}(SK, C)$ (which may be the error symbol $\perp$ if C is an invalid ciphertext).
3. $\mathcal{A}$ chooses two messages (M_0, M_1) and obtains $C^* \leftarrow \text{Encrypt}(PK, M_\gamma)$, where $\gamma \leftarrow \mathcal{U}(\{0,1\})$
4. $\mathcal{A}$ makes new decryption queries on arbitrary ciphertexts $C \neq C^*$
5. $\mathcal{A}$ outputs $\gamma' \in \{0,1\}$ and wins if $\gamma' = \gamma$

$$\mathbf{Adv}_{\mathcal{A}}^{\text{IND-CCA}}(\lambda) := \left| \Pr[\gamma' = \gamma] - \frac{1}{2} \right|$$

Remark

- In a non-adaptive chosen-ciphertext attack (CCA1), stage 4 is removed (Naor-Yung, STOC'90): no decryption query is allowed after the challenge phase
- Elgamal is not IND-CCA2-secure: $\mathcal{A}$ is given the challenge ciphertext

  $$C^* = (g^r, M_\gamma \cdot X^r) = (C_1, C_2)$$

  and can compute $C' = (C_1 \cdot g^{r'}, C_2 \cdot X^{r'}) = (g^{r+r'}, M_\gamma \cdot X^{r+r'})$, for a randomly chosen $r' \in_R \mathbb{Z}_p$, which may be submitted to the decryption oracle and reveals M_γ to $\mathcal{A}$.

0.2 Generic IND-CCA2 PKE from any IND-sID-CPA-secure IBE (Canetti-Halevi-Katz, Eurocrypt'04)

- **Keygen**(λ) : Generate $(MPK, MSK) \leftarrow \text{Setup}^{IBE}(\lambda)$.
  Choose a one-time signature scheme $\Sigma = (G, S, V)$.
  Define $PK := (MPK, \Sigma)$, $SK := MSK$.
- **Encrypt**(PK, M) :
  1. Generate a key pair $(SVK, SSK) \leftarrow G(\lambda)$ for the one-time signature.
  2. Compute $C^{IBE} \leftarrow \text{Encrypt}^{IBE}(MPK, M, SVK)$, which is an encryption of M under the identity SVK.
  3. Compute $\sigma \leftarrow S(SSK, C^{IBE})$ and output $C = (SVK, C^{IBE}, \sigma)$.
- **Decrypt**(SK, C) :
  1. Return $\perp$ if $V(SVK, C^{IBE}, \sigma) = 0$.
  2. Compute $d_{SVK} \leftarrow \text{Keygen}^{IBE}(MSK, SVK)$.
  3. Output $M \leftarrow \text{Decrypt}^{IBE}(MPK, d_{SVK}, C^{IBE})$.

Definition: Strong Unforgeability A one-time signature $\Sigma = (G, S, V)$ is strongly unforgeable under chosen-message attacks (SUF-CMA) if no PPT adversary $\mathcal{A}$ has noticeable advantage one the following game:

1. The challenger generates $(SVK, SSK) \leftarrow G(\lambda)$ and gives SVK to $\mathcal{A}$
2. $\mathcal{A}$ chooses exactly *one* message M and obtains $\sigma \leftarrow S(SSK, M)$
3. $\mathcal{A}$ outputs (M^*, σ^*) and wins if
   (a) $V(SVK, M^*, \sigma^*) = 1$
   (b) $(M^*, \sigma^*) \neq (M, \sigma)$

In many signature schemes, signatures are not unique (i.e., a given message has many valid signatures). For such schemes, the above notion is stictly stronger than the usual notion of unforgeability, where condition (b) is replaced by $M \neq M^*$.

Theorem The PKE scheme produced by the Canetti-Halevi-Katz transformation is IND-CCA2-secure assuming that

- Σ is strongly unforgeable
- The IBE scheme is IND-sID-CPA-secure

Proof Let $C^* = (SVK^*, C^{IBE^*}, \sigma^*)$ be the challenge ciphertext given to the adversary in the IND-CCA2 game. We consider two kinds of attacks:

- Type I attack: $\mathcal{A}$ makes at least one valid decryption query $C = (SVK, C^{IBE}, \sigma)$ such that $SVK \neq SVK^*$ (by "valid decryption query", we mean one where the one-time signature σ correctly verifies w.r.t. SVK).
- Type II attack: All valid decryption queries $C_i = (SVK_i, C_i^{IBE}, \sigma_i)$ contain one-time verification keys SVK_i such that $SVK_i \neq SVK^*$

Type I attack contradicts the SUF-CMA-security of Σ. The proof is straightforward and omitted here.

Let $\mathcal{A}$ be Type II adversary with noticeable advantage ε. Using $\mathcal{A}$, we build an IND-sID-CPA adversary $\mathcal{B}$ against the IBE scheme:

- $\mathcal{B}$ generates a one-time signature key pair $(SVK^*, SSK^*) \leftarrow G(\lambda)$ and declares SVK^* as its target identity $ID^* = SVK^*$ in the IND-sID-CPA security game.
- $\mathcal{B}$ obtains MPK^{IBE} from its own challenger and gives $PK = (MPK^{IBE}, \Sigma)$ to $\mathcal{A}$ as a public key for the IND-CCA security game.

  Queries: suppose that $\mathcal{A}$ queries the decryption of a ciphertext $C = (SVK, C^{IBE}, \sigma)$. Since $\mathcal{A}$ is a Type II attacker, we necessarily have $SVK \neq SVK^*$, so that $\mathcal{B}$ can obtain an IBE private key $d_{SVK} \leftarrow \text{Keygen}(MSK^{IBE}, SVK)$ from its challenger, and compute $M \leftarrow \text{Decrypt}^{IBE}(MPK^{IBE}, d_{SVK}, C)$.

  Challenge: $\mathcal{A}$ chooses (M_0, M_1) which $\mathcal{B}$ sends to its own challenger. The latter returns a challenge ciphertext $C^{IBE^*} \leftarrow \text{Encrypt}^{IBE}(MPK^{IBE}, M_\gamma, SVK^*)$ for the IND-sID-CPA game. Then, $\mathcal{B}$ computes $C^* = (SVK^*, C^{IBE^*}, \sigma^*)$ where $\sigma^* \leftarrow S(SSK^*, C^{IBE^*})$ and gives it as a challenge to $\mathcal{A}$.

Output $\mathcal{A}$ outputs $\gamma' \in \{0,1\}$ and $\mathcal{B}$ outputs γ'.

Clearly, if $\mathcal{A}$ is sucessful in the IND-CCA game, so is $\mathcal{B}$ in the IND-sID-CPA game. $\square$

Remark The CHK tranform turns any 2-level HIBE with an IND-sID-CCA2-secure IBE scheme.

0.3 Attribute-based encryption and fuzzy IBE

0.3.1 Definition

Definition: Fuzzy IBE (Sahai-Waters, Eurocrypt'05 [3])

- Decryption works when identities of ciphertext/key are close enough
- Identities are sets of descriptive attributes ("student", "EU citizen", "Driving license holder", etc)
- If a ciphertext is encrypted for an attribute set ω' and private key corresponds to attribute set ω, decryption works if $|\omega \cap \omega'| \geqslant d$ for some $d \in \mathbb{N}$.

Motivation:

- Use biometric identities (e.g., iris scan)
- Access control on encrypted data (e.g., at least 2 attributes among "research staff member", "Patent engineer", "CEO")

Selective security: Let $d \in \mathsf{poly}(\lambda)$ be the decryption threshold.

0. The adversary $\mathcal{A}$ chooses a target attribute set ω^*
1. The challenger generates $(MPK, MSK) \leftarrow \text{Setup}(\lambda, d)$ and gives MPK to $\mathcal{A}$
2. $\mathcal{A}$ makes private key queries: $\mathcal{A}$ chooses an arbitrary attribute set ω such that $|\omega \cap \omega'| < d$, and obtains $d_\omega \leftarrow \text{Keygen}(MSK, \omega)$.
3. $\mathcal{A}$ chooses (M_0, M_1) and obtains $C \leftarrow \text{Encrypt}(MPK, M_\gamma, \omega^*)$ with $\gamma \leftarrow \mathcal{U}(\{0,1\})$
4. $\mathcal{A}$ makes more private key queries
5. $\mathcal{A}$ outputs a bit $\gamma' \in \{0,1\}$ and wins if $\gamma = \gamma'$. Again, $\mathcal{A}$'s advantage is defined to be

$$\mathbf{Adv}_{\mathcal{A}}^{\text{FIBE-CPA}}(\lambda) := \left| \Pr[\gamma' = \gamma] - \frac{1}{2} \right|$$

In the stronger notion of full (a.k.a. adaptive) security, the adversary chooses ω^* at step 3 at the same time as M_0, M_1.

0.3.2 Construction for large attribute universes (Sahai-Waters, Eurocrypt'05 [3])

- **Setup**(λ, d) :
  1. Choose cyclic groups (G, G_T) of prime order $p > 2^\lambda$ with a bilinear map $e : G \times G \to G_T$ and generators $g, g_2 \in G$
  2. Choose $y \leftarrow \mathbb{Z}_p$ and computes $g_1 := g^y$
  3. Choose a function $T : \mathbb{Z}_p \to G$ (to be defined later)
     Set $MPK := ((G, G_T), g, g_1(= g^y), g_2, T)$ and $MSK := y \in \mathbb{Z}_p$

- **Keygen**(MSK, ω) : Choose a random polynomial $q(X) \in \mathbb{Z}_p[X]$ of degree $d-1$ such that $q(0) = y$. For each $i \in \omega$, choose $r_i \leftarrow \mathbb{Z}_p$ and compute $(D_i, d_i) = (g_2^{q(i)} \cdot T(i)^{r_i}, g^{r_i})$. Return the private key

  $$d_\omega = \{(D_i, d_i)\}_{i \in \omega}.$$

  Note that, for each $i \in \omega$, the pair (D_i, d_i) satisfies the relation

  $$e(D_i, g) = e(g, g_2)^{q(i)} \cdot e(T(i), d_i). \tag{1}$$

- **Encrypt**(MPK, M, ω') : To encrypt $M \in G_T$ under the attribute set ω', choose $s \leftarrow \mathbb{Z}_p$ and compute the ciphertext

  $$CT = \Big(\omega', E' = M \cdot e(g_1, g_2)^s, E = g^s, \{E_i = T(i)^s\}_{i \in \omega'}\Big).$$

- **Decrypt**(MPK, d_ω, CT) : Given $d_\omega = \{(D_i, d_i)\}_{i \in \omega}$, find a set $S \subseteq \omega \cap \omega'$ such that $|S| = d$ (or return $\perp$ if none exists). For each $i \in S$, compute

  $$\frac{e(D_i, E)}{e(E_i, d_i)} = e(g, g_2)^{q(i) \cdot s}. \tag{2}$$

  Since $e(g, g_2)^{q(0) \cdot s} = e(g_1, g_2)^s$, if we define the function

  $$\Delta_{i,S}(X) := \prod_{\substack{j \in S \\ j \neq i}} \frac{X - j}{i - j},$$

  the message M can be obtained by performing a Lagrange interpolation in the exponent and computing

  $$M = \frac{E'}{\displaystyle\prod_{i \in S} \left(\frac{e(D_i, E)}{e(d_i, E_i)}\right)^{\Delta_{i,S}(0)}}$$

The correctness of the scheme can be verified by observing that, if we raise both members of (1) to the power $s \in \mathbb{Z}_p$, we obtain (2).

Theorem The scheme provides selective security if the DBDH assumption holds.

Proof Let $\mathcal{A}$ be selective adversary with advantage ε. We build a DBDH distinguisher $\mathcal{B}$ with advantage ε. Algorithm $\mathcal{B}$ takes as input (g, g^a, g^b, g^c, Z) and uses $\mathcal{A}$ to decide if $Z = e(g,g)^{abc}$ or $Z \in_R G_T$.

The adversary $\mathcal{A}$ first chooses a target attribute set ω^*. To generate MPK, $\mathcal{B}$ defines $g_1 = g^a, g_2 = g^b$ and chooses the function $T : \mathbb{Z}_p \to G$ in such a way that $\forall x \in \mathbb{Z}_p$, we can write

$$T(x) = g_2^{F(x)} \cdot g^{J(x)},$$

for certain functions $F, J : \mathbb{Z}_p \to \mathbb{Z}_p$ (which are kept internal to $\mathcal{B}$) chosen such that

$$F(x) = 0 \ \text{ if and only if } \ x \in \omega^*.$$

The adversary $\mathcal{A}$ is given $MPK := ((G, G_T), g, g_1(= g^a), g_2(= g^b), T)$, which implicitly defines $MSK := a$ (note that MSK is not available to $\mathcal{B}$).

Queries: suppose that $\mathcal{A}$ queries a private key for ω such that $|\omega \cap \omega^*| < d$. Let $\Gamma = \omega \cap \omega^*$, and Γ' be any set such that $\Gamma \subseteq \Gamma' \subseteq \omega$, and $|\Gamma'| = d - 1$.

- For each $i \in \Gamma' \subseteq \omega^*$, chooses $\lambda_i, r_i \in \mathbb{Z}_p$, and sets

$$D_i := g_2^{\lambda_i} \cdot T(i)^{r_i}, \qquad d_i := g^{r_i}.$$

- For each $i \in \omega \backslash \Gamma'$, we know that $i \notin \omega^*$ and we thus have $T(i) = g_2^{F(i)} \cdot g^{J(i)}$ with $F(i) \neq 0$. Hence, $\mathcal{B}$ can compute

$$D' = g_2^{q(0)} \cdot T(i)^{\tilde{r}} = T(i)^r \cdot (g^a)^{-\frac{J(i)}{F(i)}} \qquad\qquad d' = g^{\tilde{r}} = g^r \cdot (g^a)^{-\frac{1}{F(i)}}$$

where $\tilde{r} = r - \frac{a}{F(i)}$ for a randomly chosen $r \in_R \mathbb{Z}_p$. In turn, this allows $\mathcal{B}$ to compute

$$D_i = D'^{\Delta_{0,s}(i)} \cdot \prod_{j \in S} g_2^{\lambda_j \Delta_{j,S}(i)} \qquad\qquad d_i = d'^{\Delta_{0,S}(i)}$$

where $S = \Gamma' \backslash \{0\}$. Then, $\mathcal{B}$ can return the complete private key

$$d_\omega = \{(D_i, d_i) = (g_2^{q(i)} \cdot T(i)^{r_i}, g^{r_i})\}_{i \in \omega}$$

to $\mathcal{A}$.

Challenge: $\mathcal{A}$ chooses two messages $M_0, M_1 \in G_T$. At this point, $\mathcal{B}$ picks $\gamma \leftarrow \{0,1\}$ and computes

$$CT^* = \left(\omega^*, E^* = M_\gamma \cdot Z, E = g^c, \{E_i = (g^c)^{J(i)}\ \}_{i \in \omega^*}\right).$$

If $Z = e(g,g)^{abc}$ then

$$CT^* = \left(\omega^*, E' = M_\gamma \cdot e(g_1, g_2)^c, E = g^c, \{E_i = T(i)^c\}_{i \in \omega^*}\right),$$

since $T(i) = g^{J(i)}$ for each $i \in \omega^*$. If $Z \in_R G_T$, we can write

$$CT^* = \left(\omega^*, E' = M_{rand} \cdot e(g_1, g_2)^c, E = g^c, \{E_i = T(i)^c\}_{i \in \omega^*}\right),$$

for some uniformly random $M_{rand} \in_R G_T$.

Output: $\mathcal{A}$ outputs a bit $\gamma' \in \{0,1\}$. Then, $\mathcal{B}$ outputs 1 (meaning that $Z = e(g,g)^{abc}$) if $\gamma' = \gamma$. Otherwise, $\mathcal{B}$ outputs 0 (meaning that $Z \in_R G_T$). It should be clear that $\mathcal{B}$'s advantage as a DBDH distinguisher is identical to $\mathcal{A}$'s advantage ε as a selective adversary.

□

In order to choose the function $T : \mathbb{Z}_p \rightarrow G$, one possibility is to fix an upper bound n on the cardinality of any attribute set ω in the scheme. The function T can be defined so as to impicitly compute a polynomial of degree n in the exponent. Namely, the master public key includes random group elements $u_0, u_1, \ldots, u_n \in_R G$ and we define $T(x) = \prod_{i=0}^{n} u_i^{(x^i)}$ for any $x \in \mathbb{Z}_p$. In the security proof, the reduction $\mathcal{B}$ can choose $F(x)$ as the polynomial $F[X] = \prod_{i \in \omega^\star}(X - i) = \sum_{i=0}^{n} f_i X^i$ and set $u_i = g_2^{f_i} \cdot g^{t_i}$, for each $i \in \{0, \ldots, n\}$, using randomly chosen $t_0, t_1, \ldots, t_n \in_R \mathbb{Z}_p$. This guarantees that $\{u_i\}_{i=0}^{n}$ have a uniform distribution.

0.4 Extension: Key-Policy Attribute-based encryption (KP-ABE)

- Ciphertext is labeled with an attribute set ω.
- Private key corresponds to an access policy P and decryption works iff $P(\omega) = 1$.

Motivation Fine-grained access control using complex policies

Example of policy P:

("Research staff" OR "Patent engineer" OR "CEO") AND ("Hired at least one year ago")

FIBE is a particular case of KP-ABE: P consists of a single gate | AND gate
OR gate
threshold gate

Bibliography

[1] C. Rackoff, D. Simon: Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack (1991)
Lecture Notes in Computer Science Volume 576, 1992, pp 433-444
Advances in Cryptology — CRYPTO '91

[2] Canetti, R., Halevi, S., Katz, J.: Chosen-Ciphertext Security from Identity-Based Encryption.
In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 207–222.
Springer, Heidelberg (2004)

[3] Sahai, A., Waters, B.: Fuzzy identity-based encryption.
In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 457–473.
Springer, Heidelberg (2005)