Advanced cryptographic primitives

Lecture 9: *Attribute-based encryption for all circuits from LWE*

Damien Stehlé
Scribe: Sebastian Scheibner

17th November 2014

Contents

0 Introduction

We will use the content of these two articles for this course:

- Fabien Laguillaumie, Adeline Langlois, and Damien Stehlé. Chiffrement avancé à partir du problème learning with errors. *Chapitre de l'ouvrage "Informatique Mathématique, une photographie en 2014"*, 2014
- Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 545–554, New York, NY, USA, 2013. ACM

1 Reminders

1.1 Learning with Error (LWE)

Distinguish with non-negligible probability ($\geq \frac{1}{n^{w(1)}}$) over $s \leftarrow U(\mathbb{Z}_q^n)$ and with non-negligible advantage ($\geq \frac{1}{2} + \frac{1}{n^{w(1)}}$) between the following distributions:

$(A, As+e)$ and (A, b) with $A \leftarrow U(\mathbb{Z}_q^{m \times n}), b \leftarrow U(\mathbb{Z}_q^m), e \leftarrow D_{\mathbb{Z}^m,\alpha,q}$, (and m may be chosen arbitrarily large)

Sampling with a trapdoor:

Let T be such that its rows form a basis of a lattice

$$\Lambda^\perp = \{x \in \mathbb{Z}^m : x^T A = 0[q]\}$$

then we can sample in polynomial time from $D_{\Lambda^\perp,\sigma,c}$ for any $\sigma \geq \max \|t_i\| \cdot \omega\left(\sqrt{\lg m}\right)$

There exists a PPT (probabilistic polynomial time) algorithm that returns (A, T) such that:

- A is within statistical distance $\leq 2^{-n}$ of $U(\mathbb{Z}_q^{m \times n})$
- T is a basis of $\Lambda^\perp$ $(T \cdot A = 0)$
- $\max \|t_i\| \leq O(m)$ (T is small)

Given T, A and $u \in \mathbb{Z}_q^n$, we can sample in PPT $x \in \mathbb{Z}^m$ such that $x^T A = u^T[q]$

$$\Pr_{x:\ x^T A = u^T}[x] \sim \exp\left(-\pi \frac{\|x\|^2}{\sigma^2}\right)$$
$$\left(\text{for } \sigma \geq \max \|t_i\| \cdot \omega\left(\sqrt{\lg m}\right)\right)$$

By LHL, we have seen that

$$\Delta\left[\left(A, x^T A\right), (A, u)\right] \leq 2^{-2n}$$

$\Rightarrow$ With probability $\geq 1 - 2^{-n}$ over $A \leftarrow U\left(\mathbb{Z}_q^{m \times n}\right)$

$$\Delta_2 = \Delta\left(x^T A, u\right) \leq 2^{-n}$$

Proof.

$$\Delta_2 = \frac{1}{2}\sum_{A,u} |Pr(D,(A,u)) - \frac{1}{q^{mn+m}}|$$

$$= \frac{1}{2}\sum_{A,u} |Pr(D,(A,-)) \cdot Pr_x(x^T A = u) - \frac{1}{q^{mn+m}}|$$

$$\geq \frac{1}{2} \sum_{\text{bad } A \text{ st. } \Delta_2 \geq 2^{-n}} \frac{1}{q^{mn}} \sum_u |Pr_x(x^T A = u) - \frac{1}{q^m}|$$

$$> \sum_{\text{bad } A} \frac{1}{q^{mn}} \cdot 2^{-n}$$

If $Pr(A\ bad) > 1^{-n}$, then $\Delta_1 > 2^{-n}$ □

For the rest: We assume we have only good A's.

Lemma 1. *The distributions of (x,u) obtained by the two following experiments are within statistical distance $2^{-\Omega(n)}$.*

Experiment 1 *Sample $x \leftarrow D_{\mathbb{Z}^m,\sigma}$, compute $u^T = x^T A[q]$*

Experiment 2 *Sample $u \leftarrow U(\mathbb{Z}_q^n)$, sample x Gaussian such that $x^T A = u^T$ with a trapdoor T*

(Already used in security proof of GPV's IBE [1], which we saw in lecture 5 of this course)

1.2 Attribute-based Encryption (ABE)

Definition 1 (Attribute-based Encryption)**.** ABE definition for circuits.

Setup $1^\lambda \rightarrow (mpk, msk)$ This algorithm takes the security parameter λ as input and returns the master public key mpk and the master secret key msk

Extract $(mpk, msk, \mathcal{C}) \rightarrow sk_\mathcal{C}$ This algorithm takes the master public key, master secret key and a circuit and returns a secret key $sk_\mathcal{C}$ associated with the circuit

Encrypt $(mpk, att, M) \rightarrow C$ This algorithm takes the master public key a vector of attributes $att = att_1, \ldots, att_l \in \{0,1\}$ and a message M and returns a ciphertext C

Decrypt $(C, sk_\mathcal{C}) \rightarrow M'$ This algorithm takes a ciphertext C and a secret key $sk_\mathcal{C}$ and returns the decrypted message M'

Correctness With probability larger than $1 - \lambda^{-\omega(1)}$ with randomness of Setup, Extract, Encrypt we have $\forall M, \forall att \in \{0,1\}^l, \forall \mathcal{C}$ such that $\mathcal{C}(att) = 1$

$$Dec(Enc(mpk, att, M), Extract(mpk, msk, \mathcal{C})) = M$$

Security

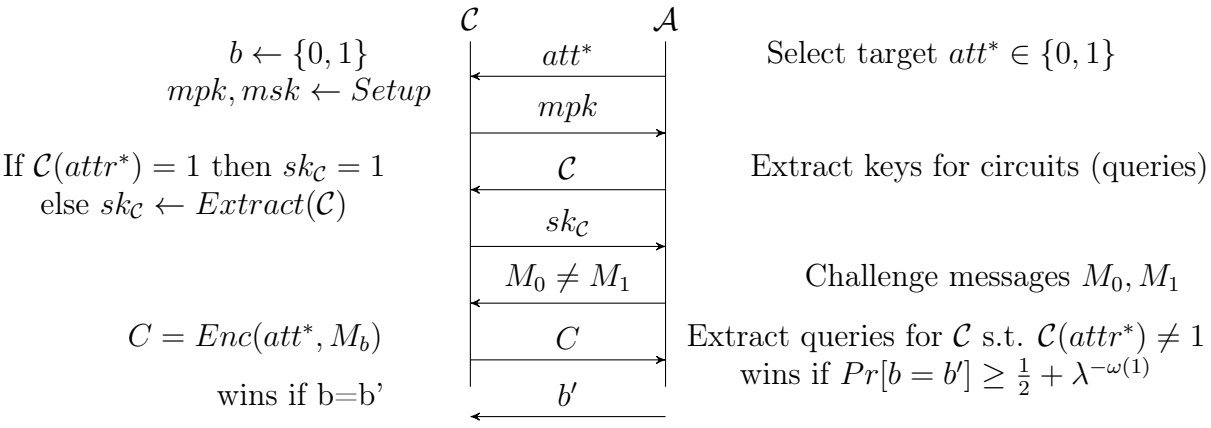

2 Two-to-One Recoding (TOR)

First we define three algorithms, where A_1, A_2, A_{tgt} are public keys for the dual-Regev encryption scheme:

Algorithm 1. *Given A_1, A_2, A_{tgt} (sampled from $U(\mathbb{Z}_q^{m \times n})$) and a trapdoor T_1, such that $T_1 A_1 = 0$*

Sample R_2 Gaussian $D_{\mathbb{Z}^{m \times m}, \sigma}$

Use T_1 to find R_1 Gaussian, such that

$$R_1 A_1 = A_{tgt} - R_2 A_2$$

(For every row x of R_1 and u of $A_{tgt} - R_2 A_2$: $x^T A_1 = u^T$)

Algorithm 2. *An algorithm using a trapdoor T_2 for A_2, it is symmetric to the first algorithm, simply replace 1 by 2 and 2 by 1.*

Algorithm 3. *Given A_1, A_2*

Sample R_1 and R_2 Gaussian (of standard parameter σ) and define

$$A_{tgt} = R_1 A_1 + R_2 A_2$$

Claim. *With probability larger than $1 - 2^{-\Omega(n)}$ over $A_1, A_2 \leftarrow U(\mathbb{Z}_q^{m \times n})$ the distributions of (A_{tgt}, R_1, R_2) obtained by running Algorithm {1,2,3} are within standard distance $2^{-\Omega(n)}$ of one another*

(follows from similar lemma on $x^T A = u^T [q]$)

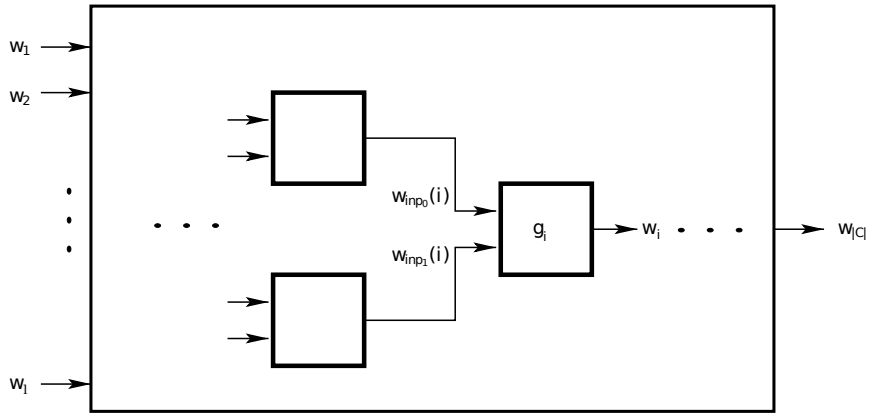

Figure 1: Description of a circuit $\mathcal{C}$ [3]

If we have $A_1, A_2, A_{tgt}, R_1, R_2$ and $b_1 = A_1 s + e1$ and $b_2 = A_2 s + e_2$ (some secret s). Then $R_1 b_1 + R_2 b_2 = A_{tgt} s + e_{tgt}$ with $e_{tgt} = R_1 e_1 + R_2 e_2$. We say s is encoded under A_{tgt} with an underlying error e_{tgt}. We have:

$$\|e_{tgt}\| \leq \text{poly}(n \lg(q)) \cdot \max(\|e_1\|, \|e_2\|)$$

$$\left(\begin{array}{l} \text{Assuming } R_1, R_2 \text{ are sampled using a trapdoor } T_1 \text{ or } T_2, \\ \sigma = \text{poly}(n \lg q) \text{ and } m \leq \text{poly}(n \lg q): \\ \\ \qquad \|\text{i'th row of } R_1\| \leq \sigma\sqrt{m} \leq \text{poly}(n \lg q) \\ \\ \text{Then use triangular inequality.} \end{array}\right)$$

3 GVW's Attribute-based Encryption (ABE)

A circuit (see Figure 1) has ℓ input wires $w_1, \ldots, w_\ell$.

Evaluation of $\mathcal{C}$ Input wires carry values in $\{0,1\}$: $\text{VAL}(w_i) \in \{0,1\}$

$$\forall i > \ell : \ \text{VAL}(w_i) = g_i(\text{VAL}(w_{inp_0}(i)), \text{VAL}(w_{inp_1}(i)))$$

output value: $\text{VAL}(w_{|\mathcal{C}|})$ the circuit accepts if $\text{VAL}(w_{|\mathcal{C}|}) = 1$

Setup The attributes $att_1, \ldots, att_\ell$ are put at the input wires. (The ℓ is fixed for all considered circuits)

For each input wire w_i, we sample $(T_0^{(i)}, A_0^{(i)}), (T_1^{(i)}, A_1^{(i)}) \in \mathbb{Z}^{m \times n} \times \mathbb{Z}_q^{m \times n}$

Sample output matrix $A^{(out)}$ uniformly, for $b \in \{0,1\}$ and $i \leq \ell$:

$$mpk = (A_{out}, (A_b^{(i)})_{b,i})$$

$$msk = (T_b^{(i)})_{b,i}$$

Enc(mpk, att, M) $M \in \{0,1\}^m$. Sample s uniformly in $\mathbb{Z}_q^n$, $e_1, \ldots, e^\ell, e_{out} \leftarrow D_{\mathbb{Z}^m, \alpha q}$

Send $\left[att, (A_{att}^{(i)} \cdot s + e_i), A^{(out)} \cdot s + e_{out} + \left[\frac{q}{2}\right] M\right]$

Extract $(\mathcal{C})$ For every internal wire w_j create $(A_0^{(j)}, T_0^{(j)}), (A_1^{(j)}, T_1^{(j)})$

Let $A_1^{(|\mathcal{C}|)} = A^{(out)}$

Sample $A_0^{(|\mathcal{C}|)}$ uniformly

For each gate $g_i : \{0,1\}^2 \to \{0,1\}$

For each $\beta_0, \beta_1 \in \{0,1\}$, create $R_0^{(i,\beta_0,\beta_1)} A_{\beta_0}^{(inp_0(i))} + R_1^{(i,\beta_0,\beta_1)} A_{\beta_1}^{(inp_1(i))} = A_{g_i(\beta_0\beta_1)}^{out(i)}$

Create $R_0^{(i,\beta_0,\beta_1)}$ and $R_1^{(i,\beta_0,\beta_1)}$ using the trapdoors $T_{\beta_0}^{(inp_0(i))}$ and $T_{\beta_1}^{(inp_1(i))}$.

$$sk_{\mathcal{C}} = (R_b^{(i,\beta_0,\beta_1)})_{i,\beta_0,\beta_1,b}$$

Dec$(C, sk_{\mathcal{C}})$

if $\mathcal{C}(att) = 0$ then stop

else for every j: Compute all VAL(w_j)

For every g_i: Take $b_{inp_0}(i)$ and $b_{inp_1}(i)$ and compute $b_{out}(i) = R_0^{(i,\beta_0,\beta_1)} \cdot b_{inp_0}(i) + R_1^{(i,\beta_0,\beta_1)} \cdot b_{inp_1}(i)$ where:

$$b_{inp_0}(i) = A_{\beta_0}^{(inp_0(i))} \cdot s + \text{"}e_0\text{"}$$
$$b_{inp_1}(i) = A_{\beta_1}^{(inp_1(i))} \cdot s + \text{"}e_1\text{"}$$
$$b_{out}(i) = A_{g_1(\beta_0,\beta_1)}^{(out(i))} \cdot s + \text{"}e_{out}\text{"}$$
$$\beta_0 = \text{VAL}(inp_0(i))$$
$$\beta_1 = \text{VAL}(inp_1(i))$$
$$\|e_{out}\| \leq \text{poly}(n \lg q) \max(\|e_0\|, \|e_1\|)$$

Correctness If $\mathcal{C}(att) = 1$:

$$A^{(out)} \cdot s + e^{out} + \left[\frac{q}{2}\right] M - (A_{\mathcal{C}(att)}^{(|\mathcal{C}|)} \cdot s + e^{|\mathcal{C}|}) = e^{out} - e^{(|\mathcal{C}|)} + \left[\frac{q}{2}\right] M$$

If $e^{out} - e^{|\mathcal{C}|}$ small, we can recover M by looking at the most significant bit of every component.

$$\|error\| \leq \text{poly}(n \lg q)^d - \|\text{fresh error}\| \leq \alpha q \, \text{poly}(n \lg q)^d$$

where d denotes the circuit depth

To decrypt correctly, we need it to be much smaller than q, which can be achieved by setting $\alpha \leq \frac{1}{\text{poly}(n \lg q)^d}$

Remark. We need to know a bound on d before setup can handle every class of depth bounded circuits.

4 Security proof

- Challenger has an LWE input A, b (uniform $As + e$), and uses the adversary to solve it.
- Adversary gives Att^* to Challenger
- Challenger defines $A^{(i)}_{att^*}$ from LWE input (it has $a \cdot b_i = A^{(i)}_{att^*} \cdot s + e_i$)
- Challenger samples $(A^{(i)}_{1-att^*_i}, T^{(i)}_{1-att^*_i})$
- Challenger defines $A^{(out)}$ from LWE input (it has $a \cdot b^{out} = A^{out} \cdot s + e^{out}$) $mpk = (A^{(i)}_b)_{b,i}, A^{(out)}$
- Extract($\mathcal{C}$) for $\mathcal{C}$ such that $\mathcal{C}(att^*) = 0$
- Challenger evaluates $\mathcal{C}$ on att^*
- For each g_i: Challenger knows
  - $A^{inp_0(i)}_{\beta^*_0}, A^{inp_1(i)}_{\beta^*_1}$ with β^*_0, β^*_1 the values taken (when evaluation in att^*)
  - $(A^{inp_0}_{1-\beta^*_0}, T^{inp_0}_{1-\beta^*_0}), (A^{inp_1}_{1-\beta^*_1}, T^{inp_1}_{1-\beta^*_1})$

  Let $\beta^* = g_i(\beta^*_0, \beta^*_1)$

  Use algorithm 3 to sample $R^{i\beta^*_0\beta^*_1}_0, R^{i\beta^*_0\beta^*_1}_1$ and define $A^{(out(i))}_{\beta^*} = R^{i\beta^*_0\beta^*_1}_0 \cdot A^{inp_0}_{\beta^*_0} + R^{i\beta^*_0\beta^*_1}_1 \cdot A^{inp_1}_{\beta^*_1}$
- Sample $A^{out}_{1-\beta^*}, T^{out}_{1-\beta^*}$

  For $(\beta_0, \beta_1) \neq (\beta^*_0, \beta^*_1)$ use $T^{inp_0}_{1-\beta^*_0}$ or $T^{inp_1}_{1-\beta^*_1}$ to create the recoding matrix as in the real scheme.

  Thanks to the correctness of TOR, the simulated $sk_{\mathcal{C}}$ has (almost) the same distribution as in the real scheme.

- Challenge phase: Challenger gets M_0, M_1 and sends back

$$(b_i)_i \text{ and } b^{out} + \left[\frac{q}{2}\right] M$$

where b_i is corresponding to Att_i^* and b_{out} is corresponding to A^{out}.

- If Challenger was given an LWE instance then

$$b_i = A^{(i)}_{att^*} s + e_i$$

$$b^{out} = A^{out} s + e^{out}$$

and this is correctly distributed ciphertext.

Otherwise $(b_i)_i, b^{out}$ are uniform, and simulated ciphertext is uniformly independent of M_0, M_1. This implies that the distinguishing advantage is 0.

References

[1] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 197–206, New York, NY, USA, 2008. ACM.

[2] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 545–554, New York, NY, USA, 2013. ACM.

[3] Fabien Laguillaumie, Adeline Langlois, and Damien Stehlé. Chiffrement avancé à partir du problème learning with errors. *Chapitre de l'ouvrage "Informatique Mathématique, une photographie en 2014"*, 2014.