Homework 3

Advanced Tools From Modern Cryptography
CS 758 : Spring 2019

Released: April 16 Saturday
Due: April 30 Sunday

FE, Lattices, Obfuscation [Total 100 pts]

1. **LWE with small secrets.** [30 pts]

   Recall that the (decision) LWE problem requires one to distinguish between the distributions of $\mathbf{r} \leftarrow \mathbb{Z}_q^m$ and $\mathbf{As} + \mathbf{e}$, where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{e} \leftarrow \chi_m$, where χ_m denotes a certain noise distribution over $\mathbb{Z}_q^m$ (for $q \geq 2$).

   Suppose you are given an algorithm D that can distinguish between the distributions of $\mathbf{r}' \leftarrow \mathbb{Z}_q^{m'}$ and $\mathbf{A}'\mathbf{s}' + \mathbf{e}'$ with a non-negligible advantage $\epsilon(n)$,[1] where $m' = m - n$, $\mathbf{A} \leftarrow \mathbb{Z}_q^{m' \times n}$, $\mathbf{s}', \mathbf{e}' \leftarrow \chi_{m'}$. Note that here $\mathbf{s}'$ is also drawn from the noise distribution, rather than the uniform distribution as in the LWE problem.

   Show that you can use the algorithm D to build a distinguisher D^* to break LWE. More precisely, D^* should have an advantage $\epsilon(n)$ of distinguishing between the distributions of $\mathbf{r} \leftarrow \mathbb{Z}_q^m$ and $\mathbf{As} + \mathbf{e}$ as in the LWE problem, but with the guarantee that $\mathbf{A}$ restricted to the first n rows is an invertible matrix (i.e., $\mathbf{A}^T = [\mathbf{A}_1^T \mid \mathbf{A}_2^T]$, where $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times n}$ is invertible).

   > This shows that LWE remains hard even when $\mathbf{s}$ is drawn from the noise distribution rather than from the uniform distribution. The condition that the first n rows $\mathbf{A}_1$ is invertible is mild: when rows of $\mathbf{A}$ are drawn uniformly randomly, one will obtain n independent rows with high probability after $O(n^2)$ samples are drawn (e.g., for a prime q, each new row is not in the linear span of prior rows with probability at least $1 - \frac{1}{q}$).
   >
   > "Modulus switching" for LWE (used in the bootstrapping of the GSW FHE scheme) relies on this.

2. **Monotone Span Programs.** [30 pts]

   A monotone access structure $\mathcal{A}$ over a groundset $[n] = \{1, \ldots, n\}$ is a subset of the power set of $[n]$[2] such that if $S \in \mathcal{A}$ and $S' \supseteq S$, then $S' \in \mathcal{A}$. We say that a pair $(\mathbf{M}, \mathbf{t})$ is a Monotone Span Program (MSP) for $\mathcal{A}$ over a field $\mathbb{F}$ if

   $$\{S \mid \exists \mathbf{v} \in \mathbb{F}^n \text{ s.t. } \mathbf{Mv} = \mathbf{t} \text{ and } \forall i \notin S, \mathbf{v}_i = 0\} = \mathcal{A}.$$

   That is, a set $S \in \mathcal{A}$ iff columns of $\mathbf{M}$ indexed by S *span* the target vector $\mathbf{t}$. Here $\mathbf{M} \in \mathbb{F}^{d \times n}$ and $\mathbf{t} \in \mathbb{F}^d$ for some integer d.

[1] Recall that an algorithm D is said to have advantage ϵ in distinguishing between two distributions X, Y if $|\Pr_{x \leftarrow X}[D(x) = 1] - \Pr_{x \leftarrow X}[D(x) = 1]| \geq \epsilon$.

[2] Power-set of a set X is the set $\{S \mid S \subseteq X\}$.

Suppose $(\mathbf{M}, \mathbf{t})$ is an MSP from some monotone access structure $\mathcal{A}$ over $[n]$, with $\mathbf{M} \in \mathbb{F}^{d \times n}$ and $\mathbf{t} \in \mathbb{F}^d \setminus \{\mathbf{0}\}$. Then, show that for any non-zero $\mathbf{t}' \in \mathbb{F}^d$ there is a matrix $\mathbf{M}' \in \mathbb{F}^{d \times n}$ such that $(\mathbf{M}', \mathbf{t}')$ is also an MSP for $\mathcal{A}$.

3. **Random Self-Reducibility of DDH** [25 points]

In this problem you need to show a worst-case to average-case reduction for the DDH problem. Let $\mathbb{G}$ be a cyclic group of prime order p, with a generator g. Suppose there is a PPT algorithm A such that

$$\Pr_{a,b \leftarrow \mathbb{Z}_p}[A(g^a, g^b, g^{ab}) = 1] < \frac{1}{2} - \epsilon \qquad \text{and} \qquad \Pr_{a,b,c \leftarrow \mathbb{Z}_p}[A(g^a, g^b, g^c) = 1] > \frac{1}{2} + \epsilon$$

Note that the probabilities are over the random choices of the input of A (as well as the randomness of A). Then, build a PPT algorithm B such that

$$\max_{a,b \in \mathbb{Z}_p} \Pr[B(g^a, g^b, g^{ab}) = 1] < \frac{1}{2} - \epsilon \qquad \text{and} \qquad \min_{\substack{a,b,c \in \mathbb{Z}_p: \\ c \neq ab}} \Pr[B(g^a, g^b, g^c) = 1] > \frac{1}{2} + \epsilon$$

where the probabilities are only over the randomness of the algorithm B. You should describe your algorithm B in terms of A, and also prove that the above property holds.

Hint: You should find a PPT transformation from (g^a, g^b, g^c) to $(g^{a'}, g^{b'}, g^{c'})$ such that for any fixed (a, b, c):

- *if $c = ab$, then (a', b') is uniformly random in $\mathbb{Z}_p^2$ and $c' = a'b'$;*
- *if $c \neq ab$, then (a', b', c') is uniformly random in $\mathbb{Z}_p^3$.*

Note that your transformation should work with a, b, c given in the exponent (so you cannot directly check if $c = ab$, unless say, $a = 0$ or $b = 0$). How much randomness will your transformation need to use for the second condition to hold?

4. **Outsourcing FHE** [5 pts]

This problem deals with reducing the encryption cost of Fully Homomorphic Encryption (FHE) for a client, who would like to outsource most of the work to an untrusted (honest-but-curious) server.

The idea is as follows: Whenever the client wishes to encrypt some data under FHE, it encrypts it under a symmetric key encryption scheme (which is a lightweight operation), and sends it to a server, who will transform it into an encryption under the FHE scheme and returns it to the client.

Describe how this idea can be implemented. In your scheme, the client may send some setup information to the server at the beginning (but afterwards, the client should only send the SKE-encrypted messages to the server). Assume that the client and server have access to the public key of the FHE.

(No proof of security required.)

5. **ABE as FE.** [5 pts]

We defined an Attribute-Based Encryption (ABE) scheme as an instance of Functional Encryption (FE) scheme with a special class of associated functions of the form

$$f_\pi(\alpha, m) = \begin{cases} (\alpha, m) & \text{if } \pi(\alpha) = 1 \\ \alpha & \text{otherwise.} \end{cases}$$

By our security definition for FE, if an adversary obtains no function keys, it should not be able to distinguish between any two messages (α_0, m_0) and (α_1, m_1) (even if $\alpha_0 \neq \alpha_1$). However, in our constructions for ABE, an encryption of (α, m) reveals α to an adversary who receives no keys.

Suggest a simple way to fix to such an ABE scheme so that it is truly a secure FE scheme for a function as defined above.

6. **Indistinguishability Obfuscation (iO)** [5 pts]

Let $\mathbb{F}$ be some family of *bijections* of the form $f : \{0,1\}^k \to \{0,1\}^k$ (k being the security parameter). Let $\mathcal{G} = \{G_{f,z} \mid f \in \mathbb{F}, z \in \{0,1\}^k\}$, where $G_{f,z} : \{0,1\}^k \to \{0,1\}^k$ is defined as

$$G_{f,z}(x) = \begin{cases} f(z) & \text{if } f(x) = 0^k \\ 0^k & \text{otherwise.} \end{cases}$$

Descibe a simple iO scheme for $\mathcal{G}$ assuming that all the functions in $\mathbb{F}$ and their inverses are efficiently computable. You should not use any computational hardness assumptions. Argue that your scheme is indeed an iO scheme. Point out where you use the efficient computability of f and f^{-1}.

Hint: What does the truth-table of $G_{f,z}$ look like? Does it have a compact representation that can be efficiently computed?

---

* **Easy KP-ABE from IBE: Proof** [Extra Credit]

We recall the construction (from lectures) of an ABE scheme for small function families, based on an IBE scheme. Let $\mathbb{F} = \{f_1, \ldots, f_t\}$ be a family of t functions (where t is at most polynomial in the security parameter), $f_i : X \to Y$. We are given an IBE scheme with ID space Y. Then we defined an ABE scheme with attribute space X and policy space[3] $\{\pi_{f,y} \mid f \in \mathbb{F}, y \in Y\}$ where $\pi_{f,y}(x) = 1$ iff $f(x) = y$, as follows:

- $\mathsf{ABE.MKeyGen}$: Output master key pair (PK, SK): $PK = \{PK_f | f \in \mathbb{F}\}$ and $SK = \{SK_f | f \in \mathbb{F}\}$, where for each $f \in \mathbb{F}$, $(PK_f, SK_f) \leftarrow \mathsf{IBE.MKeyGen}$ (generated independently).
- $\mathsf{ABE.Enc}_{PK}(m; x)$: Output (CT, x), where $CT = \{CT_f | f \in \mathbb{F}\}$ and CT_f is the IBE encryption of m for the ID $f(x)$ under the key PK_f (i.e., $CT_f = \mathsf{IBE.Enc}_{PK_f}(m; f(x))$).
- $\mathsf{ABE.FKeyGen}_{SK}(f, y)$: Output function secret-key $SK_{f,y} = \mathsf{IBE.IDKeyGen}_{SK_f}(y)$.
- $\mathsf{ABE.Dec}_{SK_{f,y}}(CT, x)$: If $f(x) = y$, output $\mathsf{IBE.Dec}_{SK_{f,y}}(CT_f)$.

Show that an adversary A with non-negligible advantage in the ABE *selective security game* can be turned into an adversary B with non-negligible advantage in the IBE selective security game. Describe how B is constructed from A, and what the advantage of B will be in terms of that of A.

You may refer to the selective security games for IBE and ABE below:

Adversary B in the IBE selective security game:

(i) outputs a challenge ID, $\widehat{ID}$;

(ii) accepts an IBE public-key PK^*; after this, at any time B can request ID-keys SK^*_{ID} for any number of IDs as long as they are different from $\widehat{ID}$;

(iii) outputs (m_0, m_1);

(iv) accepts $\mathsf{IBE.Enc}_{PK^*}(m_b; \widehat{ID})$ for $b \leftarrow \{0,1\}$, and

(v) outputs a guess for b.

Advantage:
$|\Pr[B \text{ outputs } 1 | b = 0] - \Pr[B \text{ outputs } 1 | b = 1]|$.

Adversary A in the ABE selective security game:

(i) outputs a challenge attribute $\widehat{x}$;

(ii) accepts an ABE public-key PK; after this, at any time A can request function-keys SK_π for any number of policies π as long as $\pi(\widehat{x}) = 0$;

(iii) outputs (m_0, m_1);

(iv) accepts $\mathsf{ABE.Enc}_{PK}(m_b; \widehat{x})$ for $b \leftarrow \{0,1\}$, and

(v) outputs a guess for b.

Advantage:
$|\Pr[A \text{ outputs } 1 | b = 0] - \Pr[B \text{ outputs } 1 | b = 1]|$.

Hint: You will need to use the hybrid argument: if events $H_0, \ldots, H_n$ are such that $|\Pr[H_n] - \Pr[H_0]| \geq \epsilon$, then there is some $i^* \in [n]$ such that $|\Pr[H_{i^*}] - \Pr[H_{i^*-1}]| \geq \epsilon/n$. Set up a sequence of hybrid experiments for the adversary A, and consider H_i to be the event that A outputs 1 in the i^{th} hybrid, and $|\Pr[H_n] - \Pr[H_0]|$ is A's advantage. Build B assuming that you know i^*.

[3] A policy is a function $\pi : X \to \{0,1\}$ s.t. ciphertext with attribute $x \in X$ can be decrypted using a function-key for policy π iff $\pi(x) = 1$.