Homework 3

Advanced Tools From Modern Cryptography
CS 758 : Spring 2019

Released: April 16 Tuesday
Due: April 29 Monday

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 required 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] 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. **Indistinguishability Obfuscation (iO)** [20 points]

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 use the efficient computability of f and f^{-1}.

Hint: What does the truth-table of $G_{f,z}$ look like? Can it be efficiently represented, using an efficient algorithm?

4. **ABE as FE.** [20 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). However, in our constructions for ABE, α is revealed 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.