Homework 1

Advanced Tools From Modern Cryptography
CS 758 : Autumn 2017

Released: October 29 Sunday
Due: November 10 Friday

FHE, FE, Lattices

[Total 100 pts]

1. **2-Universal Hash Function.** [20 pts]

   For a prime number q and positive integers m, n, let $D := \mathbb{Z}_q^m \setminus \{0^m\}$ and $R := \mathbb{Z}_q^n$. Below, all probabilities refer to the uniformly random choice of $\mathbf{L} \leftarrow \mathbb{Z}_q^{n \times m}$, and all addition and multiplication of numbers are modulo q.

   (a) Prove that $\forall \mathbf{x} \in D, \mathbf{a} \in R$, $\Pr_{\mathbf{L}}[\mathbf{Lx} = \mathbf{a}] = 1/|R|$.

   Hint: Fix an i s.t. $\mathbf{x}_i \neq 0$. Consider sampling $\mathbf{L}$ by picking the i^{th} column last.

   (b) Use the above to prove that $\forall \mathbf{x}, \mathbf{y} \in D$ s.t. $\mathbf{x} \neq \mathbf{y}$, $\Pr_{\mathbf{L}}[\mathbf{Lx} = \mathbf{Ly}] = 1/|R|$.

   Hint: Since $\mathbf{x} \neq \mathbf{y}$, $\mathbf{x} - \mathbf{y} \in D$.

   (c) Now suppose $D = \{0,1\}^m$. Show that $\forall \mathbf{x}, \mathbf{y} \in D$ s.t. $\mathbf{x} \neq \mathbf{y}$, $\mathbf{a}, \mathbf{b} \in R$, $\Pr_{\mathbf{L}}[\mathbf{Lx} = \mathbf{a}, \mathbf{Ly} = \mathbf{b}] = 1/|R|^2$.

   Hint: Let $\mathcal{L} = \{\mathbf{L} \mid \mathbf{Lx} = \mathbf{a}, \mathbf{Ly} = \mathbf{b}\}$. You need to argue that $|\mathcal{L}|$ does not depend on $(\mathbf{a}, \mathbf{b})$. Be explicit where all you rely on $\mathbf{x} \neq \mathbf{y}$ and that $x \in \{0,1\}^m$.

   This shows that the family of functions $\mathcal{H} = \{h_{\mathbf{L}} \mid \mathbf{L} \in \mathbb{Z}_2^{n \times m}\}$, where $h_{\mathbf{L}} : D \to R$ is defined as $h_{\mathbf{L}}(\mathbf{x}) = \mathbf{Lx}$ is a 2-universal hash function family (and has low collision probability for all prime q). We can upgrade this to a 2-universal hash function family over $D \cup \{0^m\}$ by considering $h_{\mathbf{L},\mathbf{u}}(\mathbf{x}) = \mathbf{Lx} + \mathbf{u}$ over all $(\mathbf{L}, \mathbf{u}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$

2. **LWE with small secrets.** [20 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).

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

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.

3. **Monotone Span Programs.** [20 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

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

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.

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 $\mathcal{A}$.

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.

5. **Bit OT from LWE.** [20 pts]

In the lecture we saw a passive-secure bit-OT protocol from public-key encryption (PKE) schemes in which the public-key can be sampled obliviously without knowing the secret-key. We also saw a PKE based on the hardness of LWE. Combine these two ideas to give a passive-secure bit-OT protocol. Describe the resulting OT protocol in detail (without separating out PKE as an intermediate step).

Sketch the arguments involved in the proof of security based on the hardness of LWE.

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