

1 Introduction

What if we want to process a particular file stored on a server without letting the server know any information on the file we are looking for? If we use plain public-key encryption (PKE) to encrypt those files and keep on the server, then the server has to know the secret-key which we used to encrypt our files. This is clearly insecure, since the server can access all the files which we do not want. Else, you need to download the encrypted documents that we need every time, decrypt to process them and re-encrypt to keep them back in the server. This is cumbersome and we want to avoid this. Is there a way that we can encrypt the files and store them into the server such that the server can still process any query without knowing the secret-key and send back the result we intended to get? The cryptographic primitive of “fully homomorphic” encryption (FHE) makes this possible. FHE lets us compute any arbitrary function (of our choice) on the *encrypted* data so that we would be able to use our secret key to decrypt the “encrypted output” given by the server.

In this course, we will learn some exciting techniques in modern cryptography pertaining to computing on encrypted data and outsourcing computation. We will be introduced to some beautiful primitives like FHE, Functional Encryption (FE), Indistinguishability Obfuscation (iO) and Witness Encryption (WE). To this end, we plan to cover some basic tools and techniques in this lecture for understanding FHE.

- Introduction: from PKE to FHE
- Textbook RSA (that satisfies homomorphic property with respect to only multiplication)
- An average case hard problem: Learning With Errors (LWE)
- Private-key encryption from LWE
- Public-key encryption from LWE

Notation. For an integer $n \in \mathbb{N}$, the notation $[n]$ represents the set $\{1, \dots, n\}$. We denote by $x \leftarrow \mathcal{D}$ the process of sampling a value x according to the distribution of \mathcal{D} . We consider $x \leftarrow S$ as the process of random sampling a value x according to the uniform distribution over a finite set S . We abbreviate probabilistic polynomial time as PPT. Any cryptographic scheme relies implicitly on the notion of security parameter. This parameter quantifies any PPT adversary’s probability of breaking the cryptographic protocol. We will denote security parameter by λ . We denote by $\text{negl} = \{\text{negl}_\lambda\}_{\lambda \in \mathbb{N}}$ as the class of functions over \mathbb{N} which *degrades faster than the inverse of any polynomial* in λ . In particular, $f \in \text{negl}_\lambda$ iff $\forall c \in \mathbb{N}, \exists \lambda_0 \in \mathbb{N}$ such that $\forall \lambda \geq \lambda_0, f(\lambda) \leq \lambda^{-c}$. Such a function f is called a *negligible* function.

2 Public-key Encryption

Definition 1. (Public-key encryption). A *public-key encryption* scheme for a message space \mathcal{M} is given as a tuple of PPT algorithms $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ as follows:

- $\text{PKE.Gen}(1^\lambda) \rightarrow (\text{SK}, \text{PK})$: This is a randomized key generation algorithm which takes the security parameter λ as input and outputs a public key PK and a secret key SK .
- $\text{PKE.Enc}(\text{PK}, m) \rightarrow c$: This a randomized algorithm that takes the public key PK and a message $m \in \mathcal{M}$ and outputs a ciphertext c encrypting m .
- $\text{PKE.Dec}(\text{SK}, c) \in \mathcal{M} \cup \{\perp\}$: This is a *deterministic* algorithm which takes the secret key SK and a ciphertext c as input. It outputs a message $m' \in \mathcal{M}$. (In some cases, it can also output \perp indicating a decryption failure.)

Correctness: $\forall \lambda \in \mathbb{N}, m \in \mathcal{M}$, we have

$$\Pr \left[\text{PKE.Dec}(\text{SK}, c) = m : (\text{SK}, \text{PK}) \leftarrow \text{PKE.Gen}(1^\lambda), c \leftarrow \text{PKE.Enc}(\text{PK}, m) \right] = 1$$

over the random coin tosses of the algorithms $\text{PKE.Gen}, \text{PKE.Enc}$.

Definition 2. (IND-CPA). We say that a public-key encryption scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ satisfies indistinguishability under chosen plaintext attacks (IND-CPA), if $\forall \lambda \in \mathbb{N}$ and every PPT adversary \mathcal{A} , it holds that

$$\left| \Pr \left[\text{Expt}_{\mathcal{A}}^{\text{PKE}}(1^\lambda, 0) = 1 \right] - \Pr \left[\text{Expt}_{\mathcal{A}}^{\text{PKE}}(1^\lambda, 1) = 1 \right] \right| \leq \mu(\lambda)$$

where the experiment $\text{Expt}_{\mathcal{A}}^{\text{PKE}}(1^\lambda, b)$ is defined in Figure 1 where $b \in \{0, 1\}$ and μ is a negligible function of λ .

1. The challenger obtains $\text{PKE.Gen}(1^\lambda) \rightarrow (\text{SK}, \text{PK})$ and sends PK to the adversary \mathcal{A} .
2. \mathcal{A} selects $(m_0, m_1) \in \mathcal{M}^2$ such that $|m_0| = |m_1|$ and sends (m_0, m_1) to the challenger.
3. The challenger sends $\text{PKE.Enc}(\text{PK}, m_b) \rightarrow c_b$ to \mathcal{A} .
4. The adversary \mathcal{A} outputs a guess b' for b .

Figure 1: $\text{Expt}_{\mathcal{A}}^{\text{PKE}}(1^\lambda, b)$: IND-CPA security game for public-key encryption

Note: For homomorphic encryption there is an additional (possibly PPT) algorithm Eval which takes as input the public key PK (and sometimes a separate publicly known evaluation key, Evk), a boolean function $f : \{0, 1\}^l \rightarrow \{0, 1\}$ and ciphertexts c_1, \dots, c_l . It returns a new ciphertext c^* as $\text{Eval}(\text{PK}, \text{Evk}, f, c_1, \dots, c_l) \rightarrow c^*$. Note that, Eval should be a public key algorithm.

Definition 3. (\mathcal{C} -homomorphic encryption). An encryption scheme $\text{HE} = (\text{Gen}, \text{Enc}, \text{Dec}, \text{Eval})$ (with an evaluation algorithm HE.Eval as defined above) is called a \mathcal{C} -homomorphic encryption scheme for a class of functions \mathcal{C} , a message space \mathcal{M} , if $\forall m_1, \dots, m_l \in \mathcal{M}, \forall f \in \mathcal{C}$, it holds that

$$\Pr \left[\begin{array}{l} \text{HE.Gen}(1^\lambda) \rightarrow (\text{SK}, \text{PK}, \text{Evk}), \\ \text{Dec}(\text{SK}, c^*) \neq f(m_1, \dots, m_l) : \text{HE.Enc}(\text{PK}, m_i) \rightarrow c_i, \forall i \in [l], \\ \text{Eval}(\text{PK}, \text{Evk}, f, c_1, \dots, c_l) \rightarrow c^* \end{array} \right] \in \text{negl}_\lambda.$$

over the random coin tosses of $\text{HE.Gen}, \text{HE.Enc}$.

If \mathcal{C} is a class of all polynomial sized circuits (P/poly), then we say that the scheme is *fully homomorphic*. The security requirement is the same as the IND-CPA security of the PKE.

2.1 Textbook RSA: Multiplicative homomorphic

Let us first review the textbook RSA scheme. Then we show that the RSA encryption satisfies the homomorphic property with respect to multiplication.

- $\text{Gen}(1^\lambda) \rightarrow (\text{PK}, \text{SK})$: Choose two large primes p, q and set $N = pq$. Sample $e \leftarrow \mathbb{Z}_N^*$ such that $\gcd(e, \phi(N)) = 1$. Output the public-key as $\text{PK} = (e, N)$ and secret-key as $\text{SK} = e^{-1} \bmod \phi(N)$. Here, ϕ is the Euler's totient function.
- $\text{Enc}(\text{PK}, m) \rightarrow c$: For a message $m \in \mathbb{Z}_N$, compute and return $c = m^e \bmod N$ as a ciphertext encoding m .
- $\text{Dec}(\text{SK}, c) = m'$: Compute and output message $m' = c^{\text{SK}} \bmod N$.

Homomorphic property of RSA: Let c_1 and c_2 be the RSA encryption of two messages m_1 and m_2 respectively. We show that $c_1 \cdot c_2$ produces the RSA encryption of the message $m_1 \cdot m_2 \bmod N$:

$$\begin{aligned} \text{Enc}(\text{PK}, m_1) \cdot \text{Enc}(\text{PK}, m_2) &= m_1^e \cdot m_2^e \bmod N \\ &= (m_1 \cdot m_2)^e \bmod N \\ &= \text{Enc}(\text{PK}, m_1 \cdot m_2) \end{aligned}$$

Therefore, RSA encryption naturally satisfies the multiplicative homomorphic property. Similarly, one can show that the classical El-Gamal encryption scheme is additively homomorphic. Note that RSA is not IND-CPA secure as the encryption algorithm is deterministic. There are some versions of RSA where randomization techniques (for example OAEP) are employed in the encryption to make it an IND-CPA secure scheme.

Fact: All known IND-CPA versions of RSA do not satisfy homomorphism. Hence we can note down the following problem which is open for a long time.

Prob. Construct IND-CPA multiplicative homomorphic RSA.

Therefore, it is non-trivial to construct a fully homomorphic encryption scheme from known PKEs. We discuss an average case hard problem named *Learning with Errors* (LWE) in the next section and later construct a homomorphic encryption scheme from LWE.

3 Learning With Errors

The learning with errors problem was introduced by Oded Regev [Reg09] for which he won the Gödel prize in 2018. Over this decade we have seen interesting applications of LWE in realizing many cryptographic primitives. Another advantage is that no quantum attack against LWE assumption is known till date (unlike the other problems—factoring and discrete logarithm which supplies the hardness underlying the current cryptographic schemes like RSA and ElGamal).

An learning with errors instance $\text{LWE}_{n,q,\chi}$ is parameterized by an integer $n \in \mathbb{N}$, a prime modulus q and a probability distribution χ over \mathbb{Z}_q . For a *fixed* $\mathbf{s} \in \mathbb{Z}_q^n$, we define the LWE distribution and a random distribution as follows:

- *LWE distribution*: sample $\mathbf{a} \leftarrow \mathbb{Z}_q^n, e \leftarrow \chi$ and return $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e)$
- *Random distribution*: chose $\mathbf{a} \leftarrow \mathbb{Z}_q^n, r \leftarrow \mathbb{Z}_q$ and return (\mathbf{a}, r)

For $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, let $\text{LWE}_{\mathbf{s}}$ be the oracle which outputs samples from the LWE distribution and Random be the oracle that outputs samples according to the random distribution defined above. We define two types of LWE problems below.

- Search Problem: The $\text{search-LWE}_{n,q,\chi}$ is the problem to find \mathbf{s} , given access to the oracle $\text{LWE}_{\mathbf{s}}$.
- Decision Problem: The $\text{decision-LWE}_{n,q,\chi}$ problem is to distinguish between $\text{LWE}_{\mathbf{s}}$ and Random .

The $\text{search-LWE}_{n,q,\chi}$ assumption is that the $\text{search-LWE}_{n,q,\chi}$ problem is hard and mathematically we express this as—for any PPT algorithm \mathcal{A} , it holds that

$$\Pr_{\mathbf{s} \leftarrow \mathbb{Z}_q^n} \left[\mathcal{A}^{\text{LWE}_{\mathbf{s}}}(1^n) = \mathbf{s} \right] \in \text{negl}_n$$

Similarly, the $\text{decision-LWE}_{n,q,\chi}$ assumption is that the $\text{decision-LWE}_{n,q,\chi}$ problem is hard and mathematically we represent this statement as—for any PPT distinguisher \mathcal{D} , it holds that

$$\left| \Pr [\mathcal{D}^{\text{LWE}_{\mathbf{s}}}(1^n) = 1] - \Pr [\mathcal{D}^{\text{Random}}(1^n) = 1] \right| \in \text{negl}_n$$

3.1 LWE Assumption with Fixed Number of Samples

We can redefine the above LWE assumption with an additional parameter $m \in \mathbb{N}$ which represents the number of LWE or random samples given to the adversary (instead of the access to oracles $\text{LWE}_{\mathbf{s}}$ or Rand) to solve the above problems. The m LWE samples received from $\text{LWE}_{\mathbf{s}}$ (resp., Random) are written as $(\mathbf{A}, \mathbf{A}^T \mathbf{s} + \mathbf{e})$ (resp., $(\mathbf{A}, \mathbf{A} \mathbf{r})$) where $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{e} \leftarrow \chi^m$, $\mathbf{r} \leftarrow \mathbb{Z}_q^m$. Therefore, we restate the above assumptions as follows:

Definition 4. ($\text{search-LWE}_{n,m,q,\chi}$ assumption). For any PPT algorithm \mathcal{A} , it holds that

$$\Pr_{\substack{\mathbf{s} \leftarrow \mathbb{Z}_q^n \\ \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} \\ \mathbf{e} \leftarrow \chi^m}} \left[\mathcal{A}(1^n, (\mathbf{A}, \mathbf{A}^T \mathbf{s} + \mathbf{e})) = \mathbf{s} \right] \in \text{negl}_n$$

Definition 5. ($\text{decision-LWE}_{n,m,q,\chi}$ assumption). For any PPT distinguisher \mathcal{D} and $\mathbf{r} \leftarrow \mathbb{Z}_q^m$, it holds that

$$\left| \Pr_{\substack{\mathbf{s} \leftarrow \mathbb{Z}_q^n \\ \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} \\ \mathbf{e} \leftarrow \chi^m}} \left[\mathcal{D}(1^n, (\mathbf{A}, \mathbf{A}^T \mathbf{s} + \mathbf{e})) = 1 \right] - \Pr_{\substack{\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} \\ \mathbf{r} \leftarrow \mathbb{Z}_q^m}} \left[\mathcal{D}(1^n, (\mathbf{A}, \mathbf{r})) = 1 \right] \right| \in \text{negl}_n$$

One can easily see that if the $\text{decision-LWE}_{n,m,q,\chi}$ assumption holds, then $\text{search-LWE}_{n,m,q,\chi}$ holds too. Interestingly, the opposite direction is also true with some change in parameter m . We note that the search-LWE is not suitable for cryptographic assumption as it allows an adversary to learn the secret \mathbf{s} partially (for example, the half part of \mathbf{s}). In cryptographic construction, we want that the adversary should not get any information about the secret and the decision-LWE is thus preferable.

We also define another version of LWE problem where the secret \mathbf{s} and the error \mathbf{e} are both sampled from the same distribution χ . This version of LWE is called short-secret-LWE or ss-LWE .

Claim 1. *The ss-LWE problem is as hard as normal LWE problem.*

Proof. Assume we have an oracle \mathcal{O}_{LWE} that solves (normal) LWE problem. Given $(\mathbf{A}, \mathbf{b} = \mathbf{A}^T \mathbf{s} + \mathbf{e})$ for $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \chi^n$ and $\mathbf{e} \leftarrow \chi^m$; we need to find \mathbf{b} using the above oracle. We sample $\mathbf{s}' \leftarrow \mathbb{Z}_q^n$ and define $\mathbf{b}' = \mathbf{b} + \mathbf{A}^T \mathbf{s}' = \mathbf{A}^T (\mathbf{s} + \mathbf{s}') + \mathbf{e}$. Therefore, $(\mathbf{A}, \mathbf{b}')$ becomes a normal LWE sample as $(\mathbf{s} + \mathbf{s}')$ is uniformly distributed over \mathbb{Z}_q^n . Now, $\mathcal{O}_{\text{LWE}}(\mathbf{A}, \mathbf{b}')$ outputs $\mathbf{t} \in \mathbb{Z}_q^n$ and hence we output $\mathbf{t} - \mathbf{s}'$ as the solution for the ss-LWE instance (\mathbf{A}, \mathbf{b}) .

Next, we assume that the oracle $\mathcal{O}_{\text{ss-LWE}}$ can solve an ss-LWE instance. Given a (normal) LWE instance $(\mathbf{A}, \mathbf{b} = \mathbf{A}^T \mathbf{s} + \mathbf{e})$ for $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ and $\mathbf{e} \leftarrow \chi^m$; we need to find \mathbf{b} using the above oracle. We rewrite the instance as

$$\mathbf{A}^T = \begin{pmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{pmatrix} \mathbf{s} + \begin{pmatrix} \mathbf{e}_1 \\ \mathbf{e}_2 \end{pmatrix} = \begin{pmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{pmatrix}$$

where $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times n}$, $\mathbf{A}_2 \in \mathbb{Z}_q^{(m-n) \times n}$, $\mathbf{b}_1 \in \mathbb{Z}_q^n$ and $\mathbf{b}_2 \in \mathbb{Z}_q^{m-n}$. Since \mathbf{A} is uniform over $\mathbb{Z}_q^{n \times m}$, the matrix \mathbf{A}_1 is invertible with overwhelming probability. Now, $\mathbf{b}_1 = \mathbf{A}_1 \mathbf{s} + \mathbf{e}_1$ implies that we can implicitly set $\mathbf{s} = \mathbf{A}_1^{-1}(\mathbf{b}_1 - \mathbf{e}_1)$. We rewrite \mathbf{b}_2 as

$$\begin{aligned} \mathbf{A}_2 \mathbf{s} + \mathbf{e}_2 &= \mathbf{A}_2 (\mathbf{A}_1^{-1}(\mathbf{b}_1 - \mathbf{e}_1)) + \mathbf{e}_2 \\ &= \mathbf{A}_2 \mathbf{A}_1^{-1} \mathbf{b}_1 - \mathbf{A}_2 \mathbf{A}_1^{-1} \mathbf{e}_1 + \mathbf{e}_2 \end{aligned}$$

Therefore, $\mathbf{A}_2 \mathbf{A}_1^{-1} \mathbf{b}_1 - \mathbf{b}_2 = \mathbf{A}_2 \mathbf{A}_1^{-1} \mathbf{e}_1 - \mathbf{e}_2$ which implies that $(\mathbf{A}_2 \mathbf{A}_1^{-1}, \mathbf{A}_2 \mathbf{A}_1^{-1} \mathbf{b}_1 - \mathbf{b}_2)$ is a fresh ss-LWE instance (as $\mathbf{A}_2 \mathbf{A}_1^{-1}$ is uniform over $\mathbb{Z}_q^{(m-n) \times n}$ since \mathbf{A}^T is itself uniformly chosen from $\mathbb{Z}_q^{m \times n}$ and $\mathbf{e}_1 \leftarrow \chi^n$) and hence it can be fed to $\mathcal{O}_{\text{ss-LWE}}$ to get \mathbf{e}_1 . Now, we can derive \mathbf{s} from the equation $\mathbf{s} = \mathbf{A}_1^{-1}(\mathbf{b}_1 - \mathbf{e}_1)$. \square

4 Encryption Schemes from LWE

4.1 Private-key Encryption Scheme from LWE

In this subsection we describe a private-key encryption scheme $\text{SKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ from LWE assumption. More precisely, we consider *decision-LWE* $_{n,q,\chi}$ assumption where the error distribution χ is such that $\|e\| \leq q/4$ holds with high probability for $e \leftarrow \chi$.

- $\text{SKE.Gen}(1^n)$: Sample and output $\text{SK} = \mathbf{s} \leftarrow \mathbb{Z}_q^n$.
- $\text{SKE.Enc}(\text{SK}, m)$: On input a secret-key $\text{SK} = \mathbf{s}$ and a message $m \in \{0, 1\}$, the encryption algorithm samples $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, $e \leftarrow \chi$ and output $c = (\mathbf{a}, b = (\langle \mathbf{a}, \mathbf{s} \rangle + e + m \lfloor \frac{q}{2} \rfloor) \bmod q)$.
- $\text{SKE.Dec}(\text{SK}, c)$: On input a secret-key $\text{SK} = \mathbf{s}$ and a ciphertext $c = (\mathbf{a}, b)$, the decryption algorithm first computes $u = b - \langle \mathbf{a}, \mathbf{s} \rangle$ and returns 0 if $|u| \leq q/4$; otherwise 1.

Correctness. Note that $u = b - \langle \mathbf{a}, \mathbf{s} \rangle = e + m \lfloor \frac{q}{2} \rfloor$ implies $|u| \leq q/4$ with high probability if $m = 0$ by the property of χ and $|u| > q/4$ if $m = 1$.

Security. By *decision-LWE* $_{n,q,\chi}$ assumption, $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e \bmod q)$ is computationally indistinguishable from a random vector $(\mathbf{a}, r) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$. After adding $m \lfloor \frac{q}{2} \rfloor$ to the second component the LWE sample remains indistinguishable from a random vector $(\mathbf{a}, r) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$. Therefore, no PPT adversary can distinguish between the encryption of $m = 0$ and $m = 1$.

4.2 Public-key Encryption from LWE

We describe a public-key encryption scheme $\text{PKE} = (\text{Gen}, \text{Enc}, \text{Dec})$ from LWE assumption. For the correctness of the scheme we need to chose the error distribution χ such that $\|e\| \leq \sqrt{q/4(2n+1)}$ for $e \leftarrow \chi$.

- $\text{PKE.Gen}(1^n)$: Sample $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times n}$, $\mathbf{s}, \mathbf{e} \leftarrow \chi^n$ and output $\text{SK} = \mathbf{s}$ and $\text{PK} = (\mathbf{A}, \mathbf{y} = \mathbf{A}^T \mathbf{s} + \mathbf{e}) \in \mathbb{Z}_q^{n \times n} \times \mathbb{Z}_q^n$.
- $\text{PKE.Enc}(\text{PK}, m)$: On input a public-key $\text{PK} = (\mathbf{A}, \mathbf{y})$ and a message $m \in \{0, 1\}$, the encryption algorithm samples $\mathbf{r}, \mathbf{x} \leftarrow \chi^n$, $x' \leftarrow \chi$ and output $c = (c_1, c_2) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, where $c_1 = \mathbf{A}\mathbf{r} + \mathbf{x}$, $c_2 = \langle \mathbf{y}, \mathbf{r} \rangle + x' + m \lfloor \frac{q}{2} \rfloor$.
- $\text{PKE.Dec}(\text{SK}, c)$: On input a secret-key $\text{SK} = \mathbf{s}$ and a ciphertext $c = (c_1, c_2)$, compute $u = c_2 - \langle c_1, \mathbf{s} \rangle$ and return 0, if $|u| \leq q/4$; otherwise return 1.

Correctness. To re-compute u , we first note that

$$\begin{aligned}
\langle c_1, \mathbf{s} \rangle &= \mathbf{s}^T (\mathbf{A}\mathbf{r} + \mathbf{x}) \\
&= (\mathbf{A}^T \mathbf{s})^T \mathbf{r} - \mathbf{s}^T \mathbf{x} \\
&= \langle \mathbf{A}^T \mathbf{s}, \mathbf{r} \rangle - \langle \mathbf{s}, \mathbf{x} \rangle \\
&= (\langle \mathbf{y}, \mathbf{r} \rangle - \langle \mathbf{e}, \mathbf{r} \rangle) - \langle \mathbf{s}, \mathbf{x} \rangle
\end{aligned}$$

and therefore the value of $u = c_2 - \langle c_1, \mathbf{s} \rangle$ is given by $\langle \mathbf{e}, \mathbf{r} \rangle - \langle \mathbf{s}, \mathbf{x} \rangle + x' + m \lfloor \frac{q}{2} \rfloor$. Due to the choice of χ we have $|\langle \mathbf{e}, \mathbf{r} \rangle - \langle \mathbf{s}, \mathbf{x} \rangle + x'| \leq \frac{nq}{4(2n+1)} + \frac{nq}{4(2n+1)} + \sqrt{\frac{q}{4(2n+1)}} < \frac{q}{4}$. Hence, the correctness follows.

Security. We show that the above scheme is IND-CPA secure under the harness of $\text{LWE}_{n, 2(n+1), q, \chi}$ assumption. We proceed by the following hybrid experiments.

Hybrid 0. This is the same experiment as $\text{Expt}_{\mathcal{A}}^{\text{PKE}}(1^n, 0)$ (defined in Figure 1) where the challenger encrypts with respect to bit $b = 0$.

Hybrid 1. In this experiment the adversary receives a “fake” public-key defined as $\widetilde{\text{PK}} = (\mathbf{A}, \mathbf{y}) \leftarrow \mathbb{Z}_q^{n \times n} \times \mathbb{Z}_q^n$ and ciphertexts are computed using $\widetilde{\text{PK}}$.

The indistinguishability between Hybrid 0 and Hybrid 1 is preserved by the hardness of $\text{ssLWE}_{n, n, q, \chi}$ assumption, and therefore by $\text{LWE}_{n, 2n, q, \chi}$ assumption.

Hybrid 2. Compute the challenge ciphertext as

$$\widetilde{\text{Enc}}(\widetilde{\text{PK}}, 0) = (\mathbf{a}, b)$$

where $(\mathbf{a}, b) \leftarrow \mathbb{Z}_q^n \times \mathbb{Z}_q$. Note that the challenge ciphertext in Hybrid 1 is actually $(n+1)$ ssLWE samples:

$$(\mathbf{A}\mathbf{r} + \mathbf{x}, \langle \mathbf{y}, \mathbf{r} \rangle + x')$$

where $(\mathbf{A}, \mathbf{y}) \leftarrow \mathbb{Z}_q^{n \times n} \times \mathbb{Z}_q^n$, $\mathbf{r}, \mathbf{x} \leftarrow \chi^n$, and $x' \leftarrow \chi$. Therefore, by $\text{ssLWE}_{n, n+1, q, \chi}$ or by $\text{LWE}_{n, 2n+1, q, \chi}$ assumption Hybrid 2 is computationally indistinguishable from Hybrid 1.

Hybrid 3. Compute the challenge ciphertext as

$$\widetilde{\text{Enc}}(\widetilde{\text{PK}}, 1) = (\mathbf{a}, b + \lfloor \frac{q}{2} \rfloor \bmod q)$$

where $(\mathbf{a}, b) \leftarrow \mathbb{Z}_q^n \times \mathbb{Z}_q$. Since b is chosen uniformly at random from \mathbb{Z}_q , Hybrid 2 and Hybrid 3 are indistinguishable from the adversary’s view.

Hybrid 4. We send the challenge ciphertext as

$$\text{PKE.Enc}(\widetilde{\text{PK}}, 1) = (\mathbf{A}\mathbf{r} + \mathbf{x}, \langle \mathbf{y}, \mathbf{r} \rangle + x' + \lfloor \frac{q}{2} \rfloor) \in \mathbb{Z}_q^n \times \mathbb{Z}_q.$$

Observe that Hybrid 3 and Hybrid 4 are computationally indistinguishable due to $\text{LWE}_{n,2n+1,q,\chi}$ assumption (as in Hybrid 2).

Hybrid 5. This is the same experiment as $\text{Expt}_{\mathcal{A}}^{\text{PKE}}(1^n, 1)$ (defined in Figure 1) where the challenger encrypts with respect to bit $b = 1$. Hybrid 5 is indistinguishable from Hybrid 4 by $\text{LWE}_{n,2n,q,\chi}$ assumption.

Therefore, we conclude by $\text{LWE}_{n,2(n+1),q,\chi}$ assumption that

$$|\Pr [\text{Expt}_{\mathcal{A}}^{\text{PKE}}(1^n, 0) = 1] - \Pr [\text{Expt}_{\mathcal{A}}^{\text{PKE}}(1^n, 1) = 1]| \in \text{negl}_n$$

References

- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, 2009.