CS 395T: Topics in Cryptography **Spring 2022**

Problem Set 1

Due: March 11, 2022 at 11:59pm (Submit on Gradescope) **Instructor:** David Wu

Instructions. You **must** typeset your solution in LaTeX using the provided template:

https://www.cs.utexas.edu/~dwu4/courses/sp22/static/homework.tex

You must submit your problem set via Gradescope (accessible through Canvas).

Collaboration Policy. You may discuss your general approach with other students, but you may not share written documents. You should not search online for solutions to these problems. If you do consult external sources, you must cite them in your submission. You must include the names of all of your collaborators with your submission. Refer to the official course policies for the full details.

Instructions. We will add problems to the problem set throughout the semester (roughly 1 problem each week), with the last problem added at least a week in advance of the due date. All problems are weighted equally. You should submit solutions to at least 70% of the problems (rounded down). If you submit solutions to more than 70% of the problems, we will drop the lowest-scoring problems when computing your final homework score.

Problem 1: SIS and Inhomogeneous SIS. Let n, q, β be lattice parameters where q is prime and $m, \beta = \text{poly}(n)$. Show that the $\mathsf{SIS}_{n,m+1,q,\beta}$ assumption implies the $\mathsf{ISIS}_{n,m,q,\beta'}$ assumption, where $\beta \geq \sqrt{1+(\beta')^2}$. You may treat the lattice dimension n as the security parameter. Both SIS and ISIS are defined with respect to the ℓ_2 norm.

Alternatively, show that hardness of $\mathsf{ISIS}_{n,m,q,\beta}$ implies hardness of $\mathsf{SIS}_{n,m+1,q,\beta}$.

For this problem, you only need to show **one** of the above statements. The second statement is more interesting, but a bit more challenging to prove.

Problem 2: GPV Signatures. Recall the general structure of the Gentry-Peikert-Vaikuntanathan signature scheme from lecture:

- The verification key is a random matrix $\mathbf{A} \xleftarrow{R} \mathbb{Z}_q^{n \times m}$ and the signing key is a short basis $\mathbf{B}$ for $\mathcal{L}^\perp(\mathbf{A})$.
- To sign a message $\mu \in \{0,1\}^*$, compute $\mathbf{y} \leftarrow H(\mu) \in \mathbb{Z}_q^n$, where $H\colon \{0,1\}^* \to \mathbb{Z}_q^n$ is modeled as a random oracle. Using $\mathbf{B}$, sample $\mathbf{x} \leftarrow D_{\mathcal{L}^\perp_{\mathbf{y}}(\mathbf{A}),s}$. Output the signature $\mathbf{x} \in \mathbb{Z}_q^m$.
- To verify a signature $\mathbf{x}$ on the message μ, check that $\|\mathbf{x}\| \leq \beta$ and $\mathbf{A}\mathbf{x} = H(\mu)$.

Here, n is the lattice dimension (the security parameter), $q = \text{poly}(n)$, $m = \Theta(n \log q)$, $s = \|\tilde{\mathbf{B}}\| \cdot \omega(\sqrt{\log n})$, and $\beta = s \cdot \text{poly}(\log n)$. Here $\tilde{\mathbf{B}}$ is the Gram-Schmidt basis associated with $\mathbf{B}$. You can also use the fact that $\|\tilde{\mathbf{b}}_i\| \leq \|\mathbf{b}_i\|$ for all i.

(a) Show that the above signature scheme is *insecure* as described. Specifically, sketch an efficient attack (with a high-level explanation) that breaks unforgeability of the signature scheme. An *informal* sketch is sufficient, and you may assume that there is sufficient slack in the choice of the norm bound β (as is often the case). You do *not* need to include a precise analysis of the parameters or of the adversary's advantage in your description.

(b) Describe *in one sentence* how to modify the scheme to prevent your attack above (and obtain a secure signature scheme). Justify your defense *in one sentence*. Your modification should be a simple modification to the above construction (there are multiple correct approaches).

Problem 3: Trapdoor Extension. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and let $\mathbf{R} \in \{0,1\}^{m \times k}$ such that $\mathbf{AR} = \mathbf{G}$. We showed in lecture (see also this summary) that we can use the trapdoor $\mathbf{R}$ to sample from a distribution that is statistically close to the discrete Gaussian distribution $D_{\mathcal{L}^\perp_{\mathbf{u}}(\mathbf{A}),s}$ for sufficiently large $s = \text{poly}(n, \log q)$ and any $\mathbf{u} \in \mathbb{Z}_q^n$. Take any matrix $\mathbf{A}_f \in \mathbb{Z}_q^{n \times m'}$, and let $\bar{\mathbf{A}} = [\mathbf{A} \mid \mathbf{A}_f]$. Show how to use $\mathbf{R}$ to efficiently sample from a distribution that is statistically close to $D_{\mathcal{L}^\perp_{\mathbf{u}}(\bar{\mathbf{A}}),s}$. Prove the correctness of your algorithm (and that it samples from the correct distribution). This type of trapdoor extension will be a very useful building block for constructing advanced cryptographic primitives.

Problem 4: Regev Encryption. In lecture, we described Regev encryption in the setting where the message is encoded in the *most significant bits* of the ciphertext. Here, we will consider a variant where the message is encoded in the *least significant bits* of the ciphertext. For simplicity, we will just consider the symmetric setting (but everything generalizes to the public-key setting in the manner described in lecture). Let the message space be $\mathbb{Z}_p$, n be the lattice dimension, q be the modulus, and χ be the error distribution. Suppose that $\gcd(p,q) = 1$. Note that p and q need *not* be prime here.

- The secret key is a vector $\mathbf{s} \xleftarrow{R} \mathbb{Z}_q^n$.
- To encrypt a message $\mu \in \mathbb{Z}_p$, sample $\mathbf{a} \xleftarrow{R} \mathbb{Z}_q^n$ and $e \leftarrow \chi$, and output the ciphertext $\mathsf{ct} = (\mathbf{a}, \mathbf{s}^\mathsf{T}\mathbf{a} + pe + \mu)$.

(a) Given a ciphertext ct and the secret key $\mathbf{s}$, describe the decryption algorithm. Prove correctness of the encryption scheme with your choice of decryption algorithm. You may assume that $\Pr[e \leftarrow \chi : |e| < q/(2p) - 1] = 1$.

(b) Show that under the $\mathsf{LWE}_{n,q,\chi}$ assumption, the above encryption scheme is semantically secure. (Technically, it is CPA-secure under the LWE assumption, but you do not have to show this.)

Problem 5: Key Switching and FHE. In this problem, we will develop another approach to extend Regev encryption to FHE. The advantage of this scheme over GSW is that ciphertexts are *vectors* rather than matrices. Let n be the lattice dimension, q be an *odd* modulus, and χ be a B-bounded distribution over $\mathbb{Z}_q$ (where $B = \text{poly}(n, \log q)$). Let $\mathbf{s}^\mathsf{T} = [-\tilde{\mathbf{s}}^\mathsf{T} \mid 1] \in \mathbb{Z}_q^{n+1}$ be a secret key for a Regev encryption scheme. Similar to the previous problem, we will encode the message in the *least* significant bit of the ciphertext. Namely, we say that $\mathbf{c} \in \mathbb{Z}_q^{n+1}$ is an encryption of $\mu \in \{0,1\}$ if $\mathbf{s}^\mathsf{T}\mathbf{c} = \mu + 2e$ for some small e.

(a) Let $\mathbf{t}^\mathsf{T} = [-\tilde{\mathbf{t}}^\mathsf{T} \mid 1] \in \mathbb{Z}_q^{n'+1}$ for some $n' = \text{poly}(n)$. Your goal in this problem is to construct a method that *publicly* translates a ciphertext encrypted under $\mathbf{s}$ to a ciphertext under $\mathbf{t}$. We decompose this into two algorithms:

- $\mathsf{Setup}(\mathbf{s}, \mathbf{t}) \to \mathsf{pp}$: On input $\mathbf{s} \in \mathbb{Z}_q^{n+1}$ and $\mathbf{t} \in \mathbb{Z}_q^{n'+1}$, the setup algorithm outputs a set of "key-switching" parameters pp.
- $\mathsf{KeySwitch}(\mathsf{pp}, \mathbf{c}) \to \mathbf{c}'$: On input the key-switching parameters pp and a ciphertext $\mathbf{c} \in \mathbb{Z}_q^{n+1}$, the key-switching algorithm outputs a new ciphertext $\mathbf{c}' \in \mathbb{Z}_q^{n'+1}$.

Construct efficient algorithms $(\mathsf{Setup}, \mathsf{KeySwitch})$ that satisfy the following requirements:

- If $\mathbf{c}$ is an encryption of μ with error e under secret key $\mathbf{s}$, then $\mathsf{KeySwitch}(\mathsf{pp}, \mathbf{c})$ outputs a new ciphertext $\mathbf{c}'$ that is an encryption of μ under secret key $\mathbf{t}$ with error e' where $|e'| \le |e| + \text{poly}(n, \log q)$.
- If $\mathbf{s}, \mathbf{t}$ are sampled uniformly, then the public parameters pp output by Setup are pseudorandom under the LWE assumption (i.e., they are computationally indistinguishable from uniform).

Prove that your scheme satisfies *both* of the above properties. Taken together, the above properties allows one to transform a ciphertext $\mathbf{c}$ under $\mathbf{s}$ to one under $\mathbf{t}$ without compromising semantic security of $\mathbf{c}$. **Hint:** Try setting the key-switching parameter to be an encryption of a (carefully-chosen) function of $\mathbf{s}$ under $\mathbf{t}$.

(b) Suppose you have two Regev ciphertexts $\mathbf{c}_1, \mathbf{c}_2$ encrypting μ_1 and μ_2 under $\mathbf{s} \in \mathbb{Z}_q^{n+1}$ with error magnitude at most e. Using the key-switching procedure defined above, show how to publicly and efficiently compute a Regev encryption $\mathbf{c}_\times$ of the product $\mu_1\mu_2$ under a suitably-chosen target key $\mathbf{t} \in \mathbb{Z}_q^{n+1}$. Note that $\mathbf{s}$ and $\mathbf{t}$ have

the *same* dimension. The error in $\mathbf{c}_\times$ should be bounded by $O(e^2) + \mathrm{poly}(n, \log q)$. In this setting, the public parameters for the key-switching algorithm would be included as part of the public key. Semantic security of the encryption scheme should still reduce to the LWE assumption. **Hint:** You may use the identity that for all vectors $\mathbf{u}_1, \mathbf{v}_1, \mathbf{u}_2, \mathbf{v}_2 \in \mathbb{Z}_q^n$, $(\mathbf{u}_1 \otimes \mathbf{u}_2)^\mathsf{T}(\mathbf{v}_1 \otimes \mathbf{v}_2) = (\mathbf{u}_1^\mathsf{T}\mathbf{v}_1)(\mathbf{u}_2^\mathsf{T}\mathbf{v}_2)$. Here $\mathbf{u} \otimes \mathbf{v} \in \mathbb{Z}_q^{n^2}$ denotes the vector that consists of all pairwise products $u_i v_j$ for $i, j \in [n]$.

(c) In *one* sentence, explain how you can extend the above procedure to support any (bounded) number of multiplications. You may make a circular security assumption.

Problem 6: Homomorphic Signatures from SIS. Recall the homomorphic signature scheme from lecture. The verification key vk consists of matrices $\mathbf{A}, \mathbf{B}_1, \ldots, \mathbf{B}_\ell$ and the signing key is a trapdoor td for $\mathbf{A}$. A signature on $x \in \{0,1\}^\ell$ consists of short matrices $\mathbf{R}_1, \ldots, \mathbf{R}_\ell$ where $\mathbf{A}\mathbf{R}_i = \mathbf{B}_i - x_i\mathbf{G}$ for all $i \in [\ell]$. In lecture, we showed that this scheme was (selectively) secure under the LWE assumption. In this problem, show that it also suffices to base hardness on $\mathsf{SIS}_{n,m,q,\beta}$, for some choice of $m = \Theta(n \log q)$, $\beta = m^{O(d)}$, and $q = m^{O(d)}$. Here, d is a bound on the (multiplicative) depth of the circuits supported by the homomorphic signature scheme.

Problem 7: Adaptively-Secure Homomorphic Signatures. The homomorphic signature scheme from lecture (see also the previous problem) was only proven to be selectively secure. Show how to compose the homomorphic signature scheme with a *vanilla* signature scheme to obtain a scheme that is *adaptively* secure under the *same* hardness assumption as the base homomorphic signature scheme (i.e., no complexity leveraging). Prove the adaptive single-message security of your scheme.

Optional Feedback. Please answer the following *optional* questions to help us design future problem sets. You do not need to answer these questions. However, we do encourage you to provide us feedback on how to improve the course experience.

(a) How long did you spend on this problem set?

(b) What was your favorite problem on this problem set? Why?

(c) What was your least favorite problem on this problem set? Why?

(d) Do you have any other feedback for this problem set?

(e) Do you have any other feedback on the course so far?

(f) Are there specific topics that you are interested in seeing in the second half of the course?