

Problem Set 5

Due: June 4, 2021 at 11:59pm

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

<https://crypto.stanford.edu/cs355/21sp/homework.tex>

Submission Instructions: You must submit your problem set via [Gradescope](#). Note that Gradescope requires that the solution to each problem starts on a **new page**.

Bugs: We make mistakes! If it looks like there might be a mistake in the statement of a problem, please ask a clarifying question on Piazza.

Limited Late Days: You may only use one late day on this assignment. This limitation is necessary to give us time to grade the assignment before the final grade deadline for the quarter.

Problem 1: Local Differential Privacy [10 points]. The differential-privacy model we saw in class, where a trusted curator aggregates all the data and then randomizes responses to queries, is also called the *central model* of differential privacy.

In the *local model* of differential privacy, the users do not want to trust the aggregator, so they each randomize their own data locally, before sending it to the aggregator. We'll look at a very simple local DP algorithm called *Randomized Response* (RR), which was proposed by Warner in 1965, four decades before differential privacy was invented! The goal of RR is to collect sensitive statistics (e.g., "how many people do drugs") while allowing each individual participant in the survey some amount of *deniability*.

Formally, each of the n users holds a private bit $b_i \in \{0, 1\}$. The quantity we are interested in estimating is $a := \frac{1}{n} \sum_{i=1}^n b_i$. Consider the following RR mechanism, that is run independently by each user:

- Flip two unbiased coins.
- If the first coin is heads, send b_i to the aggregator.
- Otherwise, look at the second coin:
 - If heads, send 0 to the aggregator.
 - If tails, send 1 to the aggregator.

- (a) Show that RR guarantees ϵ -differential privacy for $\epsilon = \ln(3)$ for each individual user's bit.
- (b) Let \hat{b}_i be the i -th user's randomized response. Show that the untrusted aggregator that receives all these noisy bits can compute an unbiased estimate \hat{a} of a (i.e., $\mathbb{E}[\hat{a}] = a$).
- (c) Show that the estimation error $\hat{a} - a$ has standard deviation $O(1/\sqrt{n})$.
- (d) How much worse is this than what we can achieve in the central model? Suppose all users send their bits b_i to a trusted curator that uses the Laplace mechanism to output a noisy estimate \hat{a}_c of a that is $\ln(3)$ -differentially private. Show that the estimation error $\hat{a}_c - a$ has standard deviation $O(1/n)$.

Problem 2: Private Information Retrieval [15 points]. Throughout this question, we consider one-round information-theoretic PIR over an n -bit database.

In class, we saw a simple two-server PIR with $O(n^{1/2})$ communication complexity. In this problem, you will first construct a *four*-server PIR scheme with communication complexity $O(n^{1/3})$. Then you will construct a *two*-server PIR with much improved $O(n^{1/3})$ communication complexity. As we mentioned in lecture, this $O(n^{1/3})$ scheme was essentially the best-known two-server PIR scheme for many many years, so in this problem you will reprove a very nice and very non-trivial result.

- (a) In the following box, we describe a four-server PIR scheme with $O(\sqrt{n})$ communication. Prove that the scheme is correct. Explain *informally* in 2-3 sentences why the scheme is secure as long as the adversary controls at most *one* server.

(**Hint:** Using matrix notation will make your life easy. The correctness argument should not require more than a few lines of math.)

Four-Server $O(\sqrt{n})$ -Communication PIR Scheme

Write the n -bit database as a matrix $X \in \mathbb{Z}_2^{\sqrt{n} \times \sqrt{n}}$. The client wants to read the bit X_{ij} from this database, where $i, j \in [\sqrt{n}]$. Recall that $e_i \in \mathbb{Z}_2^{\sqrt{n}}$ is the dimension- \sqrt{n} vector that is zero everywhere except with a “1” at position i .

- Query(i, j) $\rightarrow (q_{00}, q_{01}, q_{10}, q_{11})$.
 Sample random vectors $r_0, r_1, s_0, s_1 \in \mathbb{Z}_2^{\sqrt{n}}$ subject to $r_0 + r_1 = e_i \in \mathbb{Z}_2^{\sqrt{n}}$ and $s_0 + s_1 = e_j \in \mathbb{Z}_2^{\sqrt{n}}$.
 For $b_0, b_1 \in \{0, 1\}$, let $q_{b_0 b_1} \leftarrow (r_{b_0}, s_{b_1})$.
 Output $(q_{00}, q_{01}, q_{10}, q_{11})$.
- Answer(X, q) $\rightarrow a$.
 Parse the query q as a pair (r, s) with $r, s \in \mathbb{Z}_2^{\sqrt{n} \times 1}$.
 Return as the answer the single bit $a \leftarrow r^T X s \in \mathbb{Z}_2$.
- Reconstruct($a_{00}, a_{01}, a_{10}, a_{11}$) $\rightarrow X_{ij}$.
 Output $X_{ij} \leftarrow a_{00} + a_{01} + a_{10} + a_{11} \in \mathbb{Z}_2$.

- (b) Say that you have a k -server PIR scheme that requires the client to upload $U(n)$ bits to each server and download one bit from each server. Explain how to use this scheme to construct a k -server PIR scheme in which, for any $\ell \in \mathbb{N}$, each client uploads $U(n/\ell)$ bits to each server and downloads ℓ bits from each server. (You may assume that n is a multiple of ℓ .)

Sketch—without a formal proof—why your construction does not break the correctness or security of the initial PIR scheme.

- (c) Show how to combine parts (a) and (b) get a four-server PIR scheme with total communication $O(n^{1/3})$. In particular, you should calculate the optimal value of the parameter ℓ used in part (b).

- (d) Sketch how to generalize the PIR scheme in part (a) to give an eight-server PIR scheme in which the client sends $O(n^{1/3})$ bits to each server and receives a single bit from each server in return. This should only take a few sentences to describe.

(e) Now comes the grand finale! Use the *eight-server* scheme from part (d) to construct a *two-server* scheme with communication $O(n^{1/3})$.

Hints:

- Label the queries of the eight-server scheme from part-(d) as $q_{000}, q_{001}, q_{010}, \dots, q_{111}$. The two queries in your new two-server scheme should be q_{000} and q_{111} from the eight-server scheme.
- The two servers can clearly send back the 1-bit answers for q_{000} and q_{111} respectively. NOW, here is the beautiful idea: show that by sending back to the client $O(n^{1/3})$ additional bits, each of the two servers can enable the client to recover the answers for three additional queries.

Problem 3: Key-Exchange from LWE [10 points]. In this problem, we will formalize the concept of a *non-interactive key exchange* (NIKE) protocol, and then construct it from LWE. NIKE protocols are a core component of Internet protocols like TLS, and the lattice-based NIKE that we develop in this problem is a simplified variant of some of the leading candidates in the NIST competition for standardizing post-quantum key-exchange.

A *non-interactive key exchange* (NIKE) protocol for a key space \mathcal{K} consists of the following PPT algorithms:

- $\text{Setup}(1^\lambda) \rightarrow \text{pp}$: On input the security parameter λ , the setup algorithm outputs the public parameters pp .
- $\text{ClientPublish}(\text{pp}) \rightarrow (\text{priv}, \text{pub})$: On input the public parameters pp , the client-publish algorithm outputs a secret value priv , and a public message pub .
- $\text{ServerPublish}(\text{pp}) \rightarrow (\text{priv}, \text{pub})$: On input the public parameters pp , the server-publish algorithm outputs a secret value priv , and a public message pub .
- $\text{KeyGen}(\text{priv}, \text{pub}) \rightarrow \text{key}$: On input a secret value priv , and a public message pub , the key generation algorithm outputs a key $\text{key} \in \mathcal{K}$.

Correctness. We require that for all $\text{pp} \leftarrow \text{Setup}(1^\lambda)$, $(\text{pub}_0, \text{priv}_0) \leftarrow \text{ClientPublish}(\text{pp})$, $(\text{pub}_1, \text{priv}_1) \leftarrow \text{ServerPublish}(\text{pp})$, we have

$$\Pr[\text{KeyGen}(\text{priv}_0, \text{pub}_1) = \text{KeyGen}(\text{priv}_1, \text{pub}_0)] = 1 - \text{negl}(\lambda).$$

Security. For a NIKE protocol ($\text{Setup}, \text{ClientPublish}, \text{ServerPublish}, \text{KeyGen}$), we define the following two experiments:

Experiment b ($b = 0, 1$):

- The challenger computes the following:

$\text{pp} \leftarrow \text{Setup}(1^\lambda),$
 $(\text{priv}_0, \text{pub}_0) \leftarrow \text{ClientPublish}(\text{pp}),$
 $(\text{priv}_1, \text{pub}_1) \leftarrow \text{ServerPublish}(\text{pp}),$
 $\text{key}_0 \leftarrow \text{KeyGen}(\text{priv}_0, \text{pub}_1),$
 $\text{key}_1 \xleftarrow{\mathcal{R}} \mathcal{K}.$

It provides $(\text{pp}, \text{pub}_0, \text{pub}_1, \text{key}_b)$ to the adversary.

- The adversary outputs a bit $\hat{b} \in \{0, 1\}$.

Let W_b be the event that \mathcal{A} outputs 1 in Experiment b . Then, we say that a NIKE protocol is secure if

$$\left| \Pr[W_0] - \Pr[W_1] \right| = \text{negl}(\lambda).$$

- (a) Explain in words why the security definition above captures our intuitive notion of security for key-exchange.
- (b) Consider the following NIKE protocol:¹

Let $n = \text{poly}(\lambda)$, q, χ_B be parameters for which $\text{LWE}_{\text{HNF}}(n, n, q, \chi_B)$ and $\text{LWE}_{\text{HNF}}(n, n+1, q, \chi_B)$ is hard. Recall from lecture that in practice, for $\lambda = 128$, we use $n \approx 800$.

Define the key space $\mathcal{K} = \{0, 1\}$ and consider the following algorithms.

- $\text{Setup}(1^\lambda) \rightarrow \text{pp}$: Sample a matrix $\mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{n \times n}$ and set $\text{pp} = \mathbf{A}$.
- $\text{ClientPublish}(\text{pp}) \rightarrow (\text{priv}, \text{pub})$: Sample vectors $\mathbf{s} \leftarrow \chi_B^n$, $\mathbf{e} \leftarrow \chi_B^n$. Then, set $\text{priv} = \mathbf{s}$, and $\text{pub} = \mathbf{A}^T \mathbf{s} + \mathbf{e}$.
- $\text{ServerPublish}(\text{pp}) \rightarrow (\text{priv}, \text{pub})$: Sample vectors $\mathbf{s} \leftarrow \chi_B^n$, $\mathbf{e} \leftarrow \chi_B^n$. Then, set $\text{priv} = \mathbf{s}$, and $\text{pub} = \mathbf{A} \mathbf{s} + \mathbf{e}$.
- $\text{KeyGen}(\text{priv}, \text{pub}) \rightarrow \text{key}$: Let $\text{priv} = \mathbf{s} \in \mathbb{Z}_q^n$ and $\text{pub} = \mathbf{b} \in \mathbb{Z}_q^n$. The key generation algorithm first samples a small noise term $e \leftarrow \chi_B$. Then, if $\|\langle \mathbf{s}, \mathbf{b} \rangle + e\|_\infty \leq \lfloor q/4 \rfloor$, set $\text{key} = 0$. Otherwise, set $\text{key} = 1$.

Suppose that q is prime and chosen to satisfy $4nB^2/q = \text{negl}(\lambda)$. Prove that the protocol satisfies correctness. For the proof, feel free to use the following fact (you do not need to prove this fact):

For any prime q , for $\mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{n \times n}$ any two non-zero vectors $\mathbf{s}_0, \mathbf{s}_1 \in \mathbb{Z}_q^n$, and $c \in \mathbb{Z}_q$,

$$\Pr_{\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times n}} [\mathbf{s}_0^T \mathbf{A} \mathbf{s}_1 = c] = 1/q - \text{negl}(\lambda),$$

where the probability is over the random choice of \mathbf{A} .

- (c) Prove that the protocol above is secure assuming $\text{LWE}_{\text{HNF}}(n, n, q, \chi_B)$ and $\text{LWE}_{\text{HNF}}(n, n+1, q, \chi_B)$. The definition of LWE_{HNF} is on the last page of this problem set. **[Hint: Use a hybrid argument.]**

Problem 4: Time Spent [3 points for answering]. How long did you spend on this problem set? This is for calibration purposes, and the response you provide will not affect your score.

Optional Feedback [0 points]. Please answer the following questions to help us design future problem sets. You do not need to answer these questions, and if you would prefer to answer anonymously, please use this [form](#). However, we do encourage you to provide us feedback on how to improve the course experience.

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

¹We restrict the key space to $\mathcal{K} = \{0, 1\}$ for simplicity. To get a NIKE protocol for $\mathcal{K} = \{0, 1\}^{128}$, we can simply run 128 parallel instances of the protocol using the same public matrix \mathbf{A} .

- (b) What was your least favorite problem on this problem set? Why?
- (c) Do you have any other feedback for this problem set?
- (d) Do you have any other feedback on the course?

Appendix: Definition of LWE in Hermite Normal Form.

We review the formal definitions of the Learning with Errors problem in *Hermite Normal Form*. Note that in this variant of the LWE problem, the vector \mathbf{s} is sampled from the B -bounded error distribution χ_B instead of the uniform distribution. This version of the LWE problem is known to be as hard as the standard LWE problem.

$\text{LWE}_{\text{HNF}}(n, m, q, \chi_B)$: Let $n, m, q, B \in \mathbb{N}$ be positive integers, and let χ_B be a B -bounded distribution over \mathbb{Z}_q . For a given adversary \mathcal{A} , we define the following two experiments:

Experiment b ($b = 0, 1$):

- The challenger computes

$$\mathbf{A} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{m \times n}, \quad \mathbf{s} \leftarrow \chi_B^n, \quad \mathbf{e} \leftarrow \chi_B^m, \quad \mathbf{b}_0 \leftarrow \mathbf{A} \cdot \mathbf{s} + \mathbf{e}, \quad \mathbf{b}_1 \xleftarrow{\mathbb{R}} \mathbb{Z}_q^m,$$

and gives the tuple $(\mathbf{A}, \mathbf{b}_b)$ to the adversary.

- The adversary outputs a bit $\hat{b} \in \{0, 1\}$.

Let W_b be the event that \mathcal{A} outputs 1 in Experiment b . Then, we define \mathcal{A} 's advantage in solving the LWE_{HNF} problem for the set of parameters n, m, q, χ_B to be

$$\text{HNF-LWEAdv}_{n,m,q,\chi_B}[\mathcal{A}] := \left| \Pr[W_0] - \Pr[W_1] \right|.$$