

## Problem Set 3

**Problem sets are due at 4:59pm on the due date.**

Please submit your problem set, in PDF format, on Gradescope. *Each problem should be in a separate page.*

You are to work on this problem set in groups. For problem sets 1, 2, and 3, we will randomly assign the groups for the problem set. After problem set 3, you are to work on the following problem sets with groups of your choosing of size three or four. If you need help finding a group, try posting on Piazza. See the course website for our policy on collaboration. Each group member must independently write up and submit their own solutions.

*You must typeset your homework in L<sup>A</sup>T<sub>E</sub>X and submit it electronically!* Each problem answer must be provided as a separate page. Mark the top of each page with your group member names, the course number (6.5610), the problem set number and question, and the date. We have provided a template for L<sup>A</sup>T<sub>E</sub>X on the course website (see the *Psets* tab at the top of the page).

### Problem 3-1. [Extra Credit] Seminar attendance

We have another speaker coming to the security seminar this week. We will give **extra credit** on this problem set if you attend.

**Title:** Designing protocols that actually get deployed

**Speaker:** Eric Rescorla

**Abstract:** Designing good Internet protocols is hard. Designing protocols that actually get deployed is harder. We take a look at some protocols that have been widely deployed (e.g., TLS 1.3, QUIC, and the WebPKI), and some others which have been less so (e.g., IPsec, SCTP, and DNSSEC/DANE) and draw some lessons about the factors that lead to protocol success and failure.

**Bio:** Eric Rescorla has contributed extensively to many of the core security protocols used in the Internet, including TLS, DTLS, WebRTC, ACME, and QUIC. He was editor of the TLS 1.3 protocol, which secures the vast majority of Web traffic and co-founder of Let's Encrypt, a free and automated certificate authority that is now the largest on the Internet. He is the former Chief Technology Officer for Firefox and Internet Platform at Mozilla, where he was responsible for setting the overall technical strategy for the Firefox browser and Mozilla's participation in Internet standards and global policy.

Details are:

- Time: Noon on Thursday, March 7
- Place: 32-D463 (Star)
- Also important: There will be free food. We will order more than last week so that we won't run out.

*If you have a timing conflict at this time, we cannot give make-up extra credit but there will be many future extra-credit opportunities!*

There are other seminar talks as well and many are related to the content that we will cover in this class! See <https://securityseminar.csail.mit.edu/> for details.

### Problem 3-2. Random Private Information Retrieval

Consider a weaker variant of the private information retrieval problem called random private informational retrieval (RPIR), where instead of the client querying a specific database index of their choosing, they want to receive a *random* index of the database. It turns out that we can convert an RPIR scheme into a PIR scheme with higher communication costs and extra rounds of back-and-forth at the start.

For a client and server, and a database  $\mathbf{D} = (D_1, \dots, D_N) \in \{0, 1\}^N$ , define  $\text{RPIR}(1^\lambda, \mathbf{D})$ , on security parameter  $\lambda$ , to be a scheme which results in the client downloading and storing  $(j, D_j)$  where  $j$  is a random index  $j \xleftarrow{\mathbb{R}} [N]$ . Security and succinctness apply as defined in lecture. Intuitively, the server will not learn  $j$ .

We now describe how to construct a PIR scheme from RPIR. Let  $\text{PIR} = (\text{Setup}, \text{Query}, \text{Answer}, \text{Reconstruct})$  consist of the following algorithms:

- **Setup** $(1^\lambda, \mathbf{D})$ 
  - Run  $\text{RPIR}(1^\lambda, \mathbf{D})$  to produce  $(j, D_j)$  stored by client.
- **Query** $(1^\lambda, i) \rightarrow (\text{qu}, \text{st})$ 
  - Client outputs  $\text{qu} \leftarrow i \oplus j$  and client state  $\text{st} \leftarrow (i, j, D_j)$
- **Answer** $(\mathbf{D}, \text{qu}) \rightarrow \text{ans}$ 
  - Server partitions the set  $[N] = \{1, \dots, N\}$  into  $N/2$  pairs  $\{k, k \oplus \text{qu}\} = \{k, k \oplus (i \oplus j)\}$
  - For each pair  $\{k, k \oplus \text{qu}\}$ , server computes  $p_{\{k, k \oplus \text{qu}\}} = D_k \oplus D_{k \oplus \text{qu}}$ .
  - Server outputs  $\text{ans} \leftarrow \{p_{\{k, k \oplus \text{qu}\}}\}$  enumerating over  $k$  such that we have one parity in  $\text{ans}$  per pair.
- **Reconstruct** $(\text{ans}, \text{st}) \rightarrow b \in \{0, 1\}$ 
  - ???

- (a) Fill in  $\text{Reconstruct}(\text{ans}, \text{st})$ . That is, describe how the client will use its stored state and the index pair parities in  $\text{ans}$  to recover  $b = D_i$ . Assume the server can output the  $p$ 's in order so that the client knows that parity  $p_{\{k, k \oplus \text{qu}\}}$  corresponds to index  $k$  (assume the client and server can both produce the partition in the same order).
- (b) Argue in 1-2 sentences that  $\text{PIR} = (\text{Setup}, \text{Query}, \text{Answer}, \text{Reconstruct})$  satisfies security.
- (c) Argue in 1-2 sentences that  $(\text{Query}, \text{Answer})$  result in total communication  $< N$ .

### Problem 3-3. Learning parities with noise

The learning-parity-with-noise assumption is similar to the LWE assumption with two differences. First, LPN is over  $\text{GF}[2]$ ; i.e., addition and multiplication is done modulo 2. Second, the noise is a Bernoulli random variable with probability  $\epsilon$ ; i.e., each linear relation is corrupted with probability  $\epsilon$ . Specifically, the  $\text{LPN}_{n,m,\epsilon}$  assumption asserts that

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \approx (\mathbf{A}, \mathbf{u})$$

where  $\mathbf{A} \xleftarrow{\mathbb{R}} \{0, 1\}^{m \times n}$ ,  $\mathbf{s} \xleftarrow{\mathbb{R}} \{0, 1\}^n$  and  $\mathbf{e} = (e_1, \dots, e_m) \in \{0, 1\}^m$  where each  $e_i$  is a boolean random variable that is “1” with probability  $\epsilon$  and “0” with probability  $1 - \epsilon$ .

- (a) Is the  $\text{LPN}_{n,m,\epsilon}$  assumption true for  $\epsilon = 0$ ? Is it true for  $\epsilon = 1/2$ ? Is it true for  $\epsilon = 1$ ?
- (b) Assuming that the  $\text{LPN}_{n,m,\epsilon}$  assumption holds, is it true that

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}, (s_1, \dots, s_n)) \approx (\mathbf{A}, \mathbf{u}, (s_1, \dots, s_n))$$

where  $\mathbf{A} \xleftarrow{\mathbb{R}} \{0, 1\}^{m \times 2n}$ ,  $\mathbf{s} = (s_1, \dots, s_{2n}) \xleftarrow{\mathbb{R}} \{0, 1\}^{2n}$ ,  $\mathbf{u} \xleftarrow{\mathbb{R}} \{0, 1\}^m$ , and  $\mathbf{e} = (e_1, \dots, e_m) \in \{0, 1\}^m$  where each  $e_i$  is a boolean random variable that is 1 with probability  $\epsilon$ ?

- (c) **(Extra credit)** Assuming that the  $\text{LPN}_{n,m,\epsilon}$  assumption holds, is it true that

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}, s_1 \oplus \dots \oplus s_{n+1}) \approx (\mathbf{A}, \mathbf{u}, s_1 \oplus \dots \oplus s_{n+1})$$

where  $\mathbf{A} \xleftarrow{\mathbb{R}} \{0, 1\}^{(n+1) \times m}$ ,  $\mathbf{s} = (s_1, \dots, s_{n+1}) \xleftarrow{\mathbb{R}} \{0, 1\}^{n+1}$ ,  $\mathbf{u} \xleftarrow{\mathbb{R}} \{0, 1\}^m$ , and  $\mathbf{e} \in \{0, 1\}^m$  where each  $e_i$  is a boolean random variable that is 1 with probability  $\epsilon$ ?

**Problem 3-4. Fully Homomorphic Encryption**

In what follows we change the encryption algorithm of the fully homomorphic encryption scheme presented in class. We do not change its key generation algorithm or its decryption and homomorphic evaluations, so assume that we use the same parameters  $(m, n, q, \chi)$  and the same secret key  $\mathbf{s} = (-\mathbf{s}', 1)$ .

*Evaluate the following schemes on correctness, security, homomorphic addition, and homomorphic multiplication. State whether each scheme satisfies each property with a brief proof/explanation.*

(a) We change the encryption algorithm to be the following:

$\text{Enc}(\mathbf{s} \in \mathbb{Z}_q^n, b \in \{0, 1\})$ :

- Sample a random matrix  $\mathbf{A} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{(n+1) \times n}$ .
- Construct matrix  $\mathbf{B} = [\mathbf{A} \parallel \mathbf{A}\mathbf{s}']$ , where  $\mathbf{s} = (-\mathbf{s}', 1)$ .
- Output  $\mathbf{B} + \mathbf{b} \cdot \mathbf{I}_{n+1}$ , where  $\mathbf{I}_{n+1}$  is the identity matrix in  $\mathbb{Z}_q$  of dimension  $n + 1$ .

(b) We change the encryption algorithm to be the following:  $\text{Enc}(\mathbf{s} \in \mathbb{Z}_q^n, b \in \{0, 1\})$ :

- Sample a random matrix  $\mathbf{A} \xleftarrow{\mathbb{R}} \mathbb{Z}_q^{(n+1) \times n}$ .
- Sample random vector  $\mathbf{e} \xleftarrow{\mathbb{R}} \chi^{n+1}$ .
- Construct matrix  $\mathbf{B} = [\mathbf{A} \parallel \mathbf{A}\mathbf{s}' + \mathbf{e}]$ , where  $\mathbf{s} = (-\mathbf{s}', 1)$  and  $\mathbf{I}_{n+1}$  is the identity matrix in  $\mathbb{Z}_q$  of dimension  $n + 1$ .
- Output  $\mathbf{B} + \mathbf{b} \cdot \mathbf{I}_{n+1}$ .

**Problem 3-5. Linearization attack on LWE with low noise**

In this problem, you are going to implement the linearization attack on LWE. Consider the following mini example of a LWE public key (we omit the modulus for simplicity):

$$\mathbf{A} = \begin{pmatrix} 2 & 3 \\ 1 & -4 \end{pmatrix}, \mathbf{u} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} = \begin{pmatrix} 5 \\ -3 \end{pmatrix}$$

Define  $\mathbf{s} = \begin{pmatrix} x \\ y \end{pmatrix}$  and  $\mathbf{e} = \begin{pmatrix} e_0 \\ e_1 \end{pmatrix}$ , we can rewrite it as a linear system:

$$\begin{cases} 2x + 3y = 5 - e_0 \\ x - 4y = -3 - e_1 \end{cases}$$

Assuming the error is either 0 or 1, the first equation can be phrased as “ $2x + 3y$  is either 4 or 5”, so we can guarantee that the following equation is always true:

$$(2x + 3y - 4)(2x + 3y - 5) = 0$$

which is equivalent to

$$4x^2 + 12xy + 9y^2 - 18x - 27y + 20 = 0$$

If we perform the above step on each equation in the linear system, we will get  $m$  (number of rows in  $A$ ) polynomials that evaluate to 0. It is still difficult to solve systems of multivariate polynomial equations, but “linearization” comes to the rescue. The idea is simple: when we have the equation

$$4x^2 + 12xy + 9y^2 - 18x - 27y + 20 = 0$$

treat it as a linear equation with 5 variables:

$$4z_1 + 12z_2 + 9z_3 - 18z_4 - 27z_5 + 20 = 0$$

With a sufficient number of equations, we can solve the system using Gaussian elimination. In our mini example, we will have 2 equations and 5 variables, which is not enough, but if  $\mathbf{A}$  is a matrix with dimension  $5 \times 2$ , we can recover the secret key with only the public key, assuming that the equations are linearly independent. You may assume that the equations after linearization are linearly independent for the following questions.

- (a) If the dimension of  $\mathbf{A}$  is  $m \times n$  and the error is chosen in the range  $[-B, B]$  (so there are  $2B + 1$  possible error values), what's the minimum  $m$  so that the above attack works? In other words, how many variables are there after linearization?
- (b) In the file `output.txt`, you can find the information of a LWE public key. Your task is to recover the secret  $\mathbf{s}$ . Submit the value of  $\mathbf{s}$  in the pdf file and the code on Gradescope.

You can find an introduction on SageMath [here](#).