Massachusetts Institute of Technology
6.857: Network and Computer Security
Professors Ronald L. Rivest and Yael Tauman Kalai

Handout 6
April 23, 2018
**Due:** May 7, 2018

# Problem Set 5

This problem set is due on *Monday, May 7, 2018* at **11:59 PM**. Please note our late submission penalty policy in the course information handout. Please submit your problem set, in PDF format, on Gradescope. *Each problem should be in a separate PDF.* When submitting the problem in Gradescope, ensure that **all your group members are listed on Gradescope**, and not in the PDF alone.

You are to work on this problem set with groups of your choosing of size three or four. If you need help finding a group, try posting on Piazza or email `6.857-tas@mit.edu`. You don't have to tell us your group members, just make sure you indicate them on Gradescope. Be sure that all group members can explain the solutions. See Handout 1 (*Course Information*) for our policy on collaboration.

*Homework must be submitted electronically!* Each problem answer must be provided as a separate pdf. Mark the top of each page with your group member names, the course number (6.857), the problem set number and question, and the date. We have provided templates for LaTeX and Microsoft Word on the course website (see the *Resources* page).

**Grading:** All problems are worth 10 points.

With the authors' permission, we may distribute our favorite solution to each problem as the "official" solution—this is your chance to become famous! If you do not wish for your homework to be used as an official solution, or if you wish that it only be used anonymously, please note this in your profile on your homework submission.

*Our department is collecting statistics on how much time students are spending on psets, etc. For each problem, please give your estimate of the number of person-hours your team spent on that problem.*

### Problem 5-1. Auditing An Election

Suppose that you are auditing an election with 10,000 cast paper ballots. There are three candidates in this election; they are simply called candidates 1, 2, and 3. Each ballot gives a vote for one candidate. The candidate with the most votes wins.

The electronic tabulator has counted the ballots, and said that candidate 1 received the most votes. You wish to confirm that this is the correct outcome with your audit.

You have drawn a random sample of 140 ballots, and examined them by hand, obtaining counts of:

candidate 1: 60 votes

candidate 2: 50 votes

candidate 3: 30 votes

You wish to estimate "the probability (given the sample) that candidate 1 is the true winner" using a Polya's Urn computation. This runs many trials (e.g. 100,000 trials); where each trial

(a) Initializes the Urn to contain 61, 51, and 31 votes for the candidates. (The extra "+1"s correspond to the Bayesian prior.) (An "Urn" is just a fancy word for a container.)

(b) Repeats the following operation until the Urn has 10,003 votes in it: Picks a vote uniformly at random from the Urn, and add to the Urn one more vote of that type.

(c) Now the Urn has 10,003 votes in it; remove one vote for each candidate (corresponding to the three "+1" votes added in (a)). Now the Urn has 10,000 votes in it, equal to the number of cast ballots.

(d) Determines the candidate having the most votes in the Urn. (Ties broken in favor of the lower-numbered candidate, for definiteness.)

The estimated "probability that candidate 1 is the winner (given the sample)" is the fraction of such trials where candidate 1 is the winner.

**(a)** Implement this method, and estimate the probability (given the stated sample) that candidate 1 is the true winner.

The audit should escalate (enlarge the random sample) if the estimated probability is less than 0.95. Should you escalate?

**(b)** Suppose now that each ballot specified two candidates, out of a total of six eligible candidates, for election to City Council, where two of the six candidates would be elected to the Council. The two elected would be the two having the two highest number of votes. How would your audit procedure change to handle this situation?

### Problem 5-2. Outsourcing Computation to Untrusted Servers

Alice wishes to encrypt her data using FHE. However, she is worried about losing her key. She also does not want to store the key on an external server, since she is worried about the server learning her private data.

How can Alice share her secret key between two servers such that:

1. Alice can use the servers to perform computations on her encrypted data.
2. Given a ciphertext, each server can generate a "partial decryption" given his share of the secret key, such that given the two partial decryptions, Alice can efficiently generate the decrypted message.
3. Each server on its own does not learn any information either about the encrypted data or the ciphertext given for decryption (i.e., semantic security holds, even given a single share of the key).

Modify the Gentry-Sahai-Waters FHE scheme to support this scenario and argue why the scheme is secure assuming that the two servers do not communicate with each other.

### Problem 5-3. Zero Knowledge

Fix group $G$ of prime order $q$ with generator $g$. Alice has a public key $y = g^x$ and secret key $x$. She wants to prove to Bob that she knows the secret key without revealing any information about $x$.

She does this by running the following 3-round protocol with Bob:

1. Alice chooses a random $r \in \mathbb{Z}_q$ and sends $\alpha = g^r$.
2. Bob chooses a random $\beta \in \mathbb{Z}_q$ and sends $\beta$ to Alice.
3. Alice sends $\gamma = r + \beta \cdot x \bmod q$ to Bob.

Bob checks whether $g^\gamma = \alpha \cdot y^\beta$, and accepts if equality holds.

Prove the following properties about this scheme:

1. It is Honest Verifier Zero-Knowledge (HVZK).
2. It is a proof-of-knowledge protocol. Namely, argue that if there exists a cheating prover $P^*$ who convinces Bob to accept in the above protocol, one can use $P^*$ to infer $x$.