CSCI 5520: Foundations of Data Privacy
The Chinese University of Hong Kong, Spring 2015

Homework 3
due Wednesday 8 April

Please bring your solution to my office or send it over email by Wednesday 8 April. You are encouraged to collaborate on the homework and ask for assistance, but you are required to write your own solutions, list your collaborators, acknowledge any sources of help, and provide external references if you have used any.

Question 1

Closely contested elections by majority vote are very sensitive to the intention of individual voters: A small number of mistakes or miscounted votes can affect the outcome of the election. The United States presidential election in 2000 was decided by 637 votes. After George W. Bush was announced as the winner, it was found that a few thousand relevant votes were miscounted.
Consider the following mechanism *Elect* for electing a leader among Alice and Bob, who we represent by the numbers -1 and 1: Each of n voters submits a choice $x_1, \dots, x_n \in \{-1, 1\}$. The winner $w \in \{-1, 1\}$ is then chosen with probability proportional to $e^{\varepsilon w(x_1 + \cdots + x_n)}$.

(a) Show that mechanism *Elect* is dominant strategy truthful in expectation.

(b) Show that mechanism *Elect* is (4ε)-differentially private.

(c) How many more votes than Bob does Alice need in order to win with probability 99%?

Question 2

This question concerns private learning of parities from examples. A parity function is a function of the form $a(x) = \langle a, x\rangle = a_1x_1 + \cdots + a_nx_n$, where a and x are n bit strings and addition and multiplication are modulo 2. A set of examples $(x_1, y_1), \dots, (x_n, y_n)$ where $x_i \in \{0, 1\}^n$ and $y_i \in \{0, 1\}$ is *consistent* if there exists a parity $a \in \{0, 1\}^n$ such that $\langle a, x_i\rangle = y_i$ for all i.
We will analyse the following mechanism for learning parities from a database of examples.

```
Mechanism $Learn((x_1, y_1), \dots, (x_m, y_m))$:
    With probability 1/2, output $\perp$.
    Otherwise, let $S$ be a random subset of $[m]$ in which
        each index $i \in [m]$ is included independently at random with probability $\varepsilon$.
    If the examples $(x_i, y_i)\colon i \in S$ are consistent,
        Output a random $a \in \{0, 1\}^n$ such that $\langle a, x_i\rangle = y_i$ for all $i \in S$.
    Otherwise, output $\perp$.
```

(a) Let x and x' be two sets of examples that differ in their i-th entry ($(x_i, y_i) \neq (x'_i, y'_i)$). Show that for every possible output z of *Learn*,

$$\Pr[\mathit{Learn}(x) = z \mid i \in S] \leq 2\Pr[\mathit{Learn}(x') = z \mid i \notin S]$$

(**Hint:** Consider the cases of consistent and inconsistent examples separately.)

(b) Use part (a) to show that *Learn* is $\ln((1 + \varepsilon)/(1 - \varepsilon))$-differentially private.

(c) Show that if $m > 4n/\varepsilon$ and the examples are independent uniform samples of the form $(x_i, \langle a, x_i \rangle)$, $x_i \sim \{0,1\}^n$, then *Learn* outputs a with probability at least $1/4$.

(**Hint:** Lower bound the probability that a is the unique solution consistent with the examples in S: Take a union bound over all other possible solutions.)

Question 3

In this question you will show that there is no local $o(\sqrt{n}/\varepsilon)$-accurate and ε-differentially private mechanism for counting queries.

(a) Let M_1 be a local ε-differentially private algorithm over domain $\{-1, 1\}$. Show that for every possible output y_1 of M_1,

$$(1 - O(\delta\varepsilon)) \Pr[M_1(X^-) = y_1] \leq \Pr[M_1(X^+) = y_1] \leq (1 + O(\delta\varepsilon)) \Pr[M_1(X^-) = y_1]$$

where $X^+ \sim \{-1,1\}_\delta$ and $X^- \sim \{-1,1\}_{-\delta}$. (That is, $\Pr[X^+ = 1] = \Pr[X^- = -1] = (1+\delta)/2$ and $\Pr[X^+ = -1] = \Pr[X^- = 1] = (1-\delta)/2$.)

(b) Use part (a) to show that $\mathrm{Div}(M_1(X^+) \| M_1(X^-)) = O(\delta^2\varepsilon^2)$.

(c) Now let $X^+ \sim \{-1,1\}^n_\delta$, $X^- \sim \{-1,1\}^n_{-\delta}$, and $M_1, \ldots, M_n$ be local ε-differentially private algorithms. Show that

$$\mathrm{Div}\big((M_1(X_1^+), \ldots, M_n(X_n^+)) \;\|\; (M_1(X_-^1), \ldots, M_n(X_n^-))\big) = O(n\delta^2\varepsilon^2).$$

Here $M_1, \ldots, M_n$ are instantiated using independent randomness.

(d) (**Optional**) Let K be a sufficiently large constant and M be a mechanism that on input $x \in \{-1,1\}^n$ outputs a $0.1\sqrt{n}/K\varepsilon$-additive approximation to the number of 1s in x. Show that for $\delta = 1/(K\varepsilon\sqrt{n})$ and $\varepsilon \leq 1$,

$$\Pr[M(X^+) > n/2] \geq 3/4 \quad \text{and} \quad \Pr[M(X^-) > n/2] < 1/4.$$

(e) Pinsker's inequality says that for every two random variables X and Y and every event T,

$$|\Pr[X \in T] - \Pr[Y \in T]| \leq \sqrt{\tfrac{1}{2}\mathrm{Div}(X\|Y)}.$$

Use parts (c), (d), and Pinsker's inequality to conclude that there is no local, ε-differentially private, and $0.1\sqrt{n}/K\varepsilon$-accurate mechanism for counting queries.