CSCI 5520: Foundations of Data Privacy	Homework 3 Solutions
The Chinese University of Hong Kong, Spring 2015

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 $e^{\varepsilon w(x_1+\dots+x_n)}$.

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

Solution: Without loss of generality let's assume voter i prefers Bob. Let $X = x_1 + \dots + x_n$. The probability of Bob winning the election is

$$\frac{e^{\varepsilon X}}{e^{\varepsilon X} + e^{-\varepsilon X}} = \frac{1}{1 + e^{-2\varepsilon X}}.$$

For any fixing of $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n$, the term $e^{-2\varepsilon X}$ is smaller $x_i = 1$ than when $x_i = -1$, so the probability of Bob winning is larger if the i-th voter votes truthfully than if he doesn't.

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

Solution: Mechanism *Elect* is an instance of the exponential mechanism with utility $u(x, w) = w(x_1 + \dots + x_n)$ and privacy parameter 2ε. This utility function is 2-Lipschitz, so the mechanism is 4ε-differentially private.

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

Solution: The winning probability of Alice is $e^{-\varepsilon X}/(e^{\varepsilon X} - e^{-\varepsilon X}) = 1/(1 + e^{2\varepsilon X})$. This quantity exceeds 99% at the point X where $e^{2\varepsilon X}$ drops below $1/99$, or X becomes smaller than $-(\ln 99)/2\varepsilon \approx -2.30/\varepsilon$. So Alice needs about $2.30/\varepsilon$ more votes.

Question 2

This question concerns private learning of parities from examples. A parity function is a function of the form $a(x) = \langle a, x\rangle = a_1x_1 + \dots + a_nx_n$, 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.)

Solution: If the examples $(x_j, y_j), j \in S$ are inconsistent conditioned on $i \in S$, then the outcome $\perp$ occurs with probability 1 on the left hand side. Since $\perp$ occurs with probability $1/2$ regardless of the input, the inequality holds in this case.

If the examples are consistent conditioned on $i \in S$, then they remain consistent when i is taken out of S, so $\perp$ occurs with probability $1/2$ in both cases. Taking i out of S preserves all the solutions a and at most doubles the number of solutions, so the probability of producing any particular solution as output drops by at most a factor of two.

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

Solution: Using the same notation as in part (a), for any z we can write

$$\Pr[\mathit{Learn}(x) = z] = (1-\varepsilon) \Pr[\mathit{Learn}(x) = z \mid i \notin S] + \varepsilon \Pr[\mathit{Learn}(x) = z \mid i \in S]$$

The first conditional probability remains the same if we replace x by x', so using part (a) we can write

$$\begin{aligned} \Pr[\mathit{Learn}(x) = z] &\leq (1-\varepsilon) \Pr[\mathit{Learn}(x') = z \mid i \notin S] + 2\varepsilon \Pr[\mathit{Learn}(x') = z \mid i \notin S] \\ &\leq (1+\varepsilon) \Pr[\mathit{Learn}(x') = z \mid i \notin S] \\ &\leq (1+\varepsilon) \frac{\Pr[\mathit{Learn}(x') = z]}{\Pr[i \notin S]} \\ &= \frac{1+\varepsilon}{1-\varepsilon} \cdot \Pr[\mathit{Learn}(x') = z]. \end{aligned}$$

(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.)

Solution: If a' is any solution other than a, then $\langle a, x_i \rangle = \langle a', x_i \rangle$ if and only if $\langle a + a', x_i \rangle = 0$ which happens with probability exactly $1/2$ over the choice of x_i. Therefore conditioned on the choice of

S, the probability that a' is a possible output is exactly $2^{-|S|}$. By a union bound, the probability that any output other than a survives is at most $(2^n - 1)/2^{|S|}$. By the multiplicative Chernoff bound, the size of S is at most $n + 2$ with probability at least $1 - e^{-n/8}$, which is at least $3/4$ when n is sufficiently large. In this case, the probability that no solution other than a survives is at least $3/4$.

To conclude, the probability that *Learn* outputs a is at least the probability it doesn't output $\perp$ $(1/2)$ times the probability that S has size at least $n/2$ $(3/4)$ times the probability that no other solution survives conditioned on the last event $(3/4)$. This product is greater than $1/4$.

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 $\Pr[X^+ = -1] = \Pr[X^- = 1] = (1 - \delta)/2$.)

Solution: The error term in the homework statement was $O(\delta^2\varepsilon^2)$ but this was incorrect. I didn't take off points for that. The weaker bound still suffices to do the other parts.

By differential privacy, for every y_1,

$$|\Pr[M_1(1) = y_1] - \Pr[M_1(-1) = y_1]| \leq (e^\varepsilon - 1) \Pr[M_1(-1) = y_1] = O(\varepsilon).$$

Let α denote the difference between these two probabilities (without the absolute value). A short calculation shows that

$$\frac{\Pr[M_1(X^+) = y_1]}{\Pr[M_1(X^-) = y_1]} = \frac{1 + \alpha\delta}{1 - \alpha\delta}.$$

Since $|\alpha| = O(\varepsilon)$, both this ratio and its inverse are bounded by $1 + O(\varepsilon\delta)$.

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

Solution: By part (a) and using the fact that $\ln(1 + \alpha) = O(\alpha)$ we have

$$\left|\ln \frac{\Pr[M_1(X^+) = y_1]}{\Pr[M_1(X^-) = y_1]}\right| = O(\varepsilon\delta)$$

for every possible outcome y_1 of M_1. From Lemma 5 in Lecture 4 it follows that

$$\mathrm{Div}(M_1(X^+) \| M_1(X^-)) = O(\varepsilon\delta)(e^{O(\varepsilon\delta)} - 1) = O(\delta^2\varepsilon^2).$$

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

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

Here $M_1, \dots, M_n$ are instantiated using independent randomness,

Solution: By independence of the different mechanisms, the divergence in question is the sum of $\mathrm{Div}(M_i(X^+) \| M_i(X^-))$ for all i, which is at most $O(n\delta^2\varepsilon^2)$ by part (b).

(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 $\delta = 1/(\varepsilon\sqrt{n})$ and $\varepsilon \leq 1$,

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

Solution: The stronger condition $\varepsilon \leq 1/4K$ needs to be assumed for the statement to hold. The number of 1s in X^+ has mean $n/2 + \sqrt{n}/(2K\varepsilon)$ and standard deviation at most $\sqrt{n}$. For X^+ to have fewer than $n/2$ ones, the number of ones needs to deviate from its mean by at least two standard deviations. By Chebyshev's inequality, the probability of this event is less than 1/4. The other inequality is completely analogous.

(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.

Solution: Suppose M is such a mechanism. Let T be the event that the output is greater than $n/2$. If M is accurate, by part (b) the difference of probabilities is at least 1/2. By choosing K large enough we can make the divergence on the right smaller than 1/2. Pinsker's inequality then gives $1/2 < 1/2$, a contradiction.