MPRI 2.3.2, Foundations of Privacy
Final exam

The exam consists of several questions. For each of them, the percentage between parentheses indicates the percentage by which a correct answer contributes to the maximum score (20). You can answer in English or in French.

Question 1 (3 points)

Consider databases which contain at least n records of people and their age, and assume that all people have an age between 1 and 100. Consider the query "what is the average age of the people in the database, rounded to the closest natural number?". For each the the three following mechanisms to answer this query, say whether or not it is differential private. In the positive case, give the privacy parameter ε, possibly as a function of n. Justify your answer.

1. $\mathcal{K}_1(x) = 25$.
2. $\mathcal{K}_2(x) =$ Flip a fair coin. If the coin is head then the result is the average age of the people in x (rounded to the closest natural number). Otherwise, the result is 25.
3. $\mathcal{K}_3(x) =$ Flip a fair coin. If the coin is head then the result is the average age of the people in x (rounded to the closest natural number). Otherwise, the result is a number chosen randomly (i.e., with uniform probability) between 1 and 100.

Would any of the above answers change if we remove the rounding to the closest natural number ?

Question 2 (4 points)

Consider a medical database which contains records people affected by one or more of 5 diseases, **A**sthma, **B**ronchitis, **C**onjunctivitis, **D**iabetes, and **E**pilepsy. Assume that we want to study what is the most common disease, i.e. we are interested in the query "What is the disease that affects the most people in the database?". We want to define an ε-differentially private mechanism to answer this query, with a utility as high as possible. The natural notion of utility, here, is the minimization of the difference between the number of people affected by the disease y and the number of people affected by the disease y', where y is the real answer and y' is the reported answer.

One obvious mechanism would be to report the real answer with probability p, and any other possible answer with probability $pe^{-\varepsilon}$ (where p is a normalization constant). However, there is a mechanism that gives a better utility for the same level ε of privacy. Find such mechanism, prove that it is ε-differentially private, and justify the improvement in utility.

Hint: use a combination of counting queries (one for each of the diseases **A**, **B**, **C**, **D**, and **E**), add noise to them, and then compute the max.

Question 3 (3 points)

Assume that we want to make a privacy-friendly statistics on a certain population, where the values of interest are in a certain discrete set X. Assume for simplicity that each individual is associated to exactly one value $x \in X$. In order to protect the privacy, each individual will sanitize his own value using a locally-differentially private mechanism $\mathcal{K} : X \to DY$, where DY is the set of the probability distributions on Y, which is also a discrete set. Let C be the stochastic matrix induced by $\mathcal{K}$, i.e., $C_{x,y}$ is the probability that $\mathcal{K}(x)$ is y. We are interested in approximating the original distribution (expressed as a probability distribution) $p : X \to [0, 1]$. Namely, we want to compute a distribution $\hat{p}$ "as close as possible" to p. To this purpose, we collect all the results of $\mathcal{K}$ applied to each individual of the population, and then we normalize the result so to obtain a probability distribution $q : Y \to [0, 1]$. Then we apply the Iterative Bayesian Update, namely we define the following sequence of distributions $p_n : X \to [0, 1]$:

$$\begin{aligned} p_0 &= \text{uniform distribution on } X \\ p_{n+1} &= \mathcal{T}(p_n) \end{aligned}$$

where the definition of T is:

$$\mathcal{T}(r)(x) = \sum_{y \in Y} q(y) \frac{C_{xy} r(x)}{\sum_{x' \in X} C_{x'y} r(x')}$$

We stop when p_{n+1} is "close enough" to p_n, and we set the approximate distribution $\hat{p}$ to be p_{n+1}. In particular, if $p_{n+1} = p_n$ then we can stop (and set $\hat{p}$ to be p_{n+1}).

- Prove that if $\hat{p} = C^{-1}q$ then $\hat{p}$ is a fixed point of T, i.e., $\mathcal{T}(\hat{p}) = \hat{p}$.
- Prove that if p is the uniform distribution, then $\hat{p} = p$.
- Are there other cases in which $\hat{p} = p$? Justify your answer.

Question 4 (3 points)

The following is known as the "Three Prisoners problem". Three prisoners, A, B and C are sentenced to death, but one of them (uniformly chosen at random) is selected to be pardoned (so 2 out of 3 prisoners will be executed). The warden knows which one will be pardoned but is not allowed to tell the prisoners.

Prisoner A begs the warden to let him now the identity of *one of the others* who will be *executed*: "if B is pardoned, give me C's name and vice versa. If I'm pardoned, choose randomly to name B or C".

4.1 Model the problem as a channel and compute its multiplicative Bayes-capacity. Using the capacity, compute the probability of correctly guessing the pardoned prisoner after receiving the warden's answer (properly justify why the capacity is relevant for this task).

4.2 Prisoner A is of course only interested is finding information *about himself*, i.e. whether *he* is going to be executed or not. Is the warden's answer useful for A? Justify your answer by an Information Flow analysis of a properly constructed channel (either the same channel of Question 4.1 or a different one).

Question 5 (3 points)

Let C^n denote the n repeated independent runs channel.

5.1 Show that for any n, C is non-interfering iff C^n is non-interfering.

5.2 Let D be a deterministic channel. Show that, for any n, D and D^n have exactly the same leakage (for all priors and gain functions) in three different ways, using

1. Hyper-distributions
2. Partition refinement
3. Composition refinement

Question 6 (4 points)

6.1 Consider two programs on *uniformly distributed 64-bit integers*

- C_1: if $x \,\%\, 8 == 0$ then $y = x$ else $y = 0$
  completely reveals x one-eighth of the time
- C_2: $y = x \mid 00 \ldots 0111$
  always reveals all but the last three bits of x

1 Show that both channels have the same Bayes-leakage

2 Consider an adversary that has 3 tries to guess the secret. Model such an adversary by a suitable gain function and show that C_2 leaks more than C_1 under that gain function.

3 Does C_2 leaks no less than C_1 for all priors ang gain functions?

6.2 Show that in Crowds, increasing the probability of forwarding always leads to a safer protocol.

As always, properly justify all answers.