Week 11 Exercises (ECE 598 DA)

Exercise (Two-Party Sum with Differential Privacy): Two companies, A and B, want to compute the total number of customers they have (combined), without revealing their individual customer counts to each other. Let a be A's customer count and b be B's count. They want the result to satisfy ϵ-DP so that neither can infer the other's count too precisely from the output. Describe a simple two-party protocol to solve the problem.

Exercise (Randomized Response): Suppose each party in a group of n uses the randomized response mechanism (with probability 1/2 to report truthfully and 1/2 to report a random bit) to respond to a sensitive yes/no question, as described earlier. Show that this mechanism satisfies ϵ-differential privacy for an appropriate ϵ, and determine that ϵ.

Exercise (Privacy Proof for Laplace Mechanism in Multi-Party Setting): Consider a protocol where n parties use a secure aggregation to compute the exact sum S of their inputs, and then one designated party adds Laplace noise $\mathsf{Lap}(0, b)$ with scale $b = \frac{\Delta}{\epsilon}$ (where Δ is the sensitivity of the sum function) to S and publishes the result Y. Formally argue that this protocol is ϵ-differentially private for each party's input.

Exercise (Combining MPC and DP): Suppose we have an MPC protocol that can compute *any function* exactly with no leakage (except the output). If we want to implement a differentially private function via this MPC, what steps should we take? Specifically, how can we use such an MPC to answer a database query with differential privacy?

Exercise (Secure AND): Two parties, Alice and Bob, each have a private input bit (x and y, respectively). They want to securely compute the AND of their bits (i.e., output $z = x \wedge y$ to both) in the *semi-honest model*. Outline a simple protocol for this task and explain why it is secure.

Exercise (CDP vs. DP Equivalence): Prove that in the central model, any mechanism that is $(\varepsilon, 0)$-IND-CDP must also be $(\varepsilon, 0)$-DP.

Exercise (PRG-Based Mechanism): Consider a counting query $q(\mathbf{x}) = \sum_i x_i$ on a database $\mathbf{x} = (x_1, \ldots, x_n)$ with $x_i \in \{0, 1\}$. Define $M_s(\mathbf{x}) = q(\mathbf{x}) + G(s) \cdot L$, where $s \leftarrow \{0, 1\}^k$ is a random seed, $G : \{0, 1\}^k \to \mathbb{R}$ outputs a pseudorandom "noise" drawn (say) from a discrete Laplace distribution, and $L > 0$ is a scale. Show that if G is a secure PRG, then M_s is ε-CDP for ε-approximate counting (by choosing L suitably).

Exercise (Simulation Implies Indistinguishability): Let M_κ be $(\varepsilon, 0)$-SIM-CDP via simulator S_κ (where S_κ is an ε-DP mechanism). Show that M_κ is also $(\varepsilon, \text{neg}(\kappa))$-IND-CDP.