Week 6 Exercises (ECE 598 DA)

Exercise (Randomized Response under LDP): Suppose each user has a bit $x_i \in \{0,1\}$ and applies randomized response: with probability $p = \frac{e^\varepsilon}{1+e^\varepsilon}$ they report their true bit, and with probability $1-p$ they flip it.

1. Show that this mechanism satisfies ε-local differential privacy.
2. Suppose n users participate. Derive an unbiased estimator of the true mean $\mu = \frac{1}{n}\sum_i x_i$ from the randomized reports, and compute its variance.

Exercise (Shuffling Amplification): Consider n users each applying an ε_0-LDP randomizer $R_i(x_i)$. The outputs are passed through a uniformly random shuffler before the server sees them.

1. State an informal privacy amplification by shuffling theorem.
2. Give a proof sketch (at a high level) of why shuffling amplifies privacy.

Exercise (Sensitivity in Secure Aggregation): Suppose n users each have data $x_i \in [0,1]$, and a secure aggregation protocol computes the sum $S = \sum_i x_i$ with Laplace noise $\mathsf{Lap}(0, 1/\varepsilon)$ added.

1. Argue that the global sensitivity of S with respect to changing one x_i is 1.
2. Show that the noised aggregate $\tilde{S} = S + \mathsf{Lap}(0, 1/\varepsilon)$ is $(\varepsilon, 0)$-DP.
3. Derive a $(1-\beta)$ high-probability additive error bound for $\tilde{S}$.

Exercise (Federated Learning with DP-SGD): In DP-SGD, each client's per-example gradient is clipped to norm C before adding Gaussian noise.

1. Explain formally why clipping is necessary for bounding sensitivity of the (average) gradient.
2. Consider a single round without subsampling (all n clients used). Let the per-example clipped gradients be $\bar{g}_i$ with $\|\bar{g}_i\|_2 \le C$. The server releases

$$\widetilde{g} \;=\; \frac{1}{n}\sum_{i=1}^{n} \bar{g}_i \;+\; \mathcal{N}(0, \sigma^2 C^2 I).$$

   Prove that this is $(\varepsilon_0, \delta_0)$-DP (central model) by calibrating σ using the Gaussian mechanism, and give the resulting high-probability error.

3. Briefly discuss how *Poisson subsampling* with sampling rate q changes the per-round privacy (amplification by sampling).