Differential Privacy Mechanism for `Prio3` and `PureDpDiscreteLaplace`

David Cook, dcook@divviup.org

July 18, 2024

Recall the definitions of pure differential privacy and the discrete Laplace distribution from [1], and the definition of global sensitivity from [2].

Definition 1. A randomized algorithm $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies ε-differential privacy if, for all neighboring datasets $x, x' \in \mathcal{X}^n$ (differing on a single element), and all events $E \subseteq \mathcal{Y}$, we have $\mathbb{P}[M(x) \in E] \leq e^{\varepsilon} \cdot \mathbb{P}[M(x') \in E]$.

Definition 2. The discrete Laplace distribution, with scale parameter t, is defined by the following probability density function, supported on the integers.

$$\forall x \in \mathbb{Z}, \underset{X \leftarrow \text{Lap}_{\mathbb{Z}}(t)}{\mathbb{P}}[X = x] = \frac{e^{1/t} - 1}{e^{1/t} + 1} \cdot e^{-|x|/t}$$

Definition 3. The global sensitivity of a query function $f(x)$ on a dataset x is the maximum distance between two query outputs over any neighboring datasets. Here, we will use the ℓ_1 metric to measure distances between results, and the replacement definition of neighboring datasets.

$$GS_f = \max_{x,x' : \text{neighboring}} \|f(x) - f(x')\|_1$$

The following differential privacy mechanism is implemented for the combination of the `PureDpDiscreteLaplace` strategy and the `Prio3Histogram` or `Prio3SumVec` VDAFs. Let $f(x)$ be the VDAF's aggregation function, operating over the integers. The aggregation function produces a query result $q = f(x) \in \mathcal{Y}$. Without loss of generality, we assume the domain $\mathcal{Y}$ is a vector of integers, $\mathcal{Y} = \mathbb{Z}^d$. Let $\mathbb{F}_p$ be field of prime order over which `Prio3` operates. Noise is sampled from the discrete Laplace distribution $\text{Lap}_{\mathbb{Z}}(GS_f/\varepsilon)$, projected into the field, and added to each coordinate of aggregate share field element vectors. Let $\pi_{\mathbb{F}_p} : \mathbb{Z} \to \mathbb{F}_p$ and $\pi_{\mathbb{Z}} : \mathbb{F}_p \to \mathbb{Z}$ be the natural projections between the integers and field elements, where $\pi_{\mathbb{Z}}$ maps field elements to $[0, p)$. Let $\vec{\pi}_{\mathbb{F}_p} : \mathbb{Z}^d \to \mathbb{F}_p^d$ be the natural extension to project vectors of integers into vectors of field elements. Let $\vec{q}^* = f^*(x) \in \mathbb{F}_p^d$ be the element-wise projections of $\vec{q}$ and f into the field using $\vec{\pi}_{\mathbb{F}_p}$. The un-noised aggregate shares produced by `Prio3` are secret shares of the query result, $\vec{q}^* = \vec{q}^{(0)} + \vec{q}^{(1)}$. Each aggregator

samples noise from the discrete Laplace distribution and adds it to the un-noised aggregate shares, and then sends the sum as their aggregate share to the collector. If we pessimistically assume that only one honest aggregator out of the two aggregators is adding differential privacy noise, then the mechanism produces $\vec{M}(x) = \vec{q}^{(0)} + \vec{q}^{(1)} + \vec{\pi}_{\mathbb{F}_p}\left(\vec{Z}\right) = \vec{q}^{*} + \vec{\pi}_{\mathbb{F}_p}\left(\vec{Z}\right)$, where $Z_j \leftarrow \mathrm{Lap}_{\mathbb{Z}}\left(GS_f/\varepsilon\right)$ is drawn independently for all $1 \leq j \leq d$.

Theorem 4. *$\vec{M}(x) = \vec{\pi}_{\mathbb{F}_p}(f(x)) + \vec{\pi}_{\mathbb{F}_p}\left(\vec{Z}\right), Z_j \leftarrow Lap_{\mathbb{Z}}\left(GS_f/\varepsilon\right)$ satisfies ε-differential privacy.*

Proof. We will show Definition 1 holds for singleton events, where E is a set of cardinality one, then other events will follow by a union bound.

For neighboring datasets x and x', let $\vec{q} = f(x)$, $\vec{q}' = f(x')$, and $\vec{q}^{*} = \vec{\pi}_{\mathbb{F}_p}(f(x))$, and let q_j, q_j^{*}, and Z_j denote the j-th component of the respective vectors. Then $M_j(x) = q_j^{*} + \pi_{\mathbb{F}_p}(Z_j)$. Applying the probability density function of the discrete Laplace distribution, we have:

$$\forall j \in [d], y_j \in \mathbb{F}_p, \mathbb{P}\left[M_j(x) = y_j\right] = \mathbb{P}\left[\pi_{\mathbb{F}_p}(Z_j) = y_j - q_j^{*}\right]$$

$$= \sum_{k=-\infty}^{\infty} \mathbb{P}\left[Z_j = \pi_{\mathbb{Z}}(y_j) - q_j + kp\right]$$

$$= \sum_{k=-\infty}^{\infty} \frac{e^{\varepsilon/GS_f} - 1}{e^{\varepsilon/GS_f} + 1} \exp\left(\frac{-\varepsilon\left|\pi_{\mathbb{Z}}(y_j) - q_j + kp\right|}{GS_f}\right)$$

$$= \frac{e^{\varepsilon/GS_f} - 1}{e^{\varepsilon/GS_f} + 1} \sum_{k=-\infty}^{\infty} \exp\left(\frac{-\varepsilon\left|\pi_{\mathbb{Z}}(y_j) - q_j + kp\right|}{GS_f}\right)$$

Since each Z_j is drawn independently, the probability of the mechanism returning some result can be found by taking the product of the probabilities for each dimension of the result vector.

$$\mathbb{P}\left[\vec{M}(x) = \vec{y}\right] = \left(\frac{e^{\varepsilon/GS_f} - 1}{e^{\varepsilon/GS_f} + 1}\right)^{d} \prod_{j=1}^{d} \sum_{k=-\infty}^{\infty} \exp\left(\frac{-\varepsilon\left|\pi_{\mathbb{Z}}(y_j) - q_j + kp\right|}{GS_f}\right)$$

$$\mathbb{P}\left[\vec{M}(x') = \vec{y}\right] = \left(\frac{e^{\varepsilon/GS_f} - 1}{e^{\varepsilon/GS_f} + 1}\right)^{d} \prod_{j=1}^{d} \sum_{k=-\infty}^{\infty} \exp\left(\frac{-\varepsilon\left|\pi_{\mathbb{Z}}(y_j) - q_j' + kp\right|}{GS_f}\right)$$

By the definition of global sensitivity, we know $\|\vec{q} - \vec{q}'\|_{\ell_1} \leq GS_f$. We can break up the ℓ_1 distance between $\vec{q}$ and $\vec{q}'$ by dimension, and relate this sum of absolute values of differences to the product of multiplicative factors of $e^{|q_j - q_j'|}$, in order to get the bound we need. Let $\delta_j = q_j - q_j'$. By the triangle inequality, $\left|\pi_{\mathbb{Z}}(y_j) - q_j' + kp\right| \leq \left|\pi_{\mathbb{Z}}(y_j) - q_j + kp\right| + |\delta_j|$. Since $\varepsilon > 0$ and $GS_f > 0$, then,

$$-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}(y_j) - q_j' + kp\right| \geq -\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}(y_j) - q_j + kp\right| - \frac{\varepsilon}{GS_f}\left|\delta_j\right|$$

$$\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j'+kp\right|\right)\geq\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j+kp\right|-\frac{\varepsilon\left|\delta_j\right|}{GS_f}\right)$$

$$\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j'+kp\right|\right)\geq\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j+kp\right|\right)e^{-\frac{\varepsilon\left|\delta_j\right|}{GS_f}}$$

$$e^{\frac{\varepsilon\left|\delta_j\right|}{GS_f}}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j'+kp\right|\right)\geq\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j+kp\right|\right)$$

$$\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j+kp\right|\right)\leq e^{\frac{\varepsilon\left|\delta_j\right|}{GS_f}}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j'+kp\right|\right)$$

Since the above holds for a fixed y, q and q', and any j and k, we can first add and then multiply inequalities together.

$$\sum_{k=-\infty}^{\infty}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j+kp\right|\right)\leq e^{\frac{\varepsilon\left|\delta_j\right|}{GS_f}}\sum_{k=-\infty}^{\infty}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j'+kp\right|\right)$$

$$\prod_{j=1}^{d}\sum_{k=-\infty}^{\infty}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j+kp\right|\right)\leq\prod_{j=1}^{d}e^{\frac{\varepsilon\left|\delta_j\right|}{GS_f}}\sum_{k=-\infty}^{\infty}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j'+kp\right|\right)$$

$$\prod_{j=1}^{d}\sum_{k=-\infty}^{\infty}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j+kp\right|\right)\leq\exp\left(\frac{\varepsilon\sum_{j=1}^{d}\left|\delta_j\right|}{GS_f}\right)\prod_{j=1}^{d}\sum_{k=-\infty}^{\infty}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j'+kp\right|\right)$$

$$\prod_{j=1}^{d}\sum_{k=-\infty}^{\infty}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j+kp\right|\right)\leq\exp\left(\frac{\varepsilon}{GS_f}\left\|\vec{q}-\vec{q}'\right\|_{\ell_1}\right)\prod_{j=1}^{d}\sum_{k=-\infty}^{\infty}\exp\left(-\frac{\varepsilon}{GS_f}\left|\pi_{\mathbb{Z}}\left(y_j\right)-q_j'+kp\right|\right)$$

Then, since $\|\vec{q} - \vec{q}'\|_{\ell_1} \leq GS_f$,

$$\prod_{j=1}^{d} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{\varepsilon}{GS_f} \left|\pi_{\mathbb{Z}}\left(y_j\right) - q_j + kp\right|\right) \leq$$
$$e^{\frac{\varepsilon}{GS_f} GS_f} \prod_{j=1}^{d} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{\varepsilon}{GS_f} \left|\pi_{\mathbb{Z}}\left(y_j\right) - q'_j + kp\right|\right)$$

$$\prod_{j=1}^{d} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{\varepsilon}{GS_f} \left|\pi_{\mathbb{Z}}\left(y_j\right) - q_j + kp\right|\right) \leq$$
$$e^{\varepsilon} \prod_{j=1}^{d} \sum_{k=-\infty}^{\infty} \exp\left(-\frac{\varepsilon}{GS_f} \left|\pi_{\mathbb{Z}}\left(y_j\right) - q'_j + kp\right|\right)$$

This shows that, for any neighboring x and x', and any y, $\mathbb{P}\left[\vec{M}\left(x\right) = \vec{y}\right] \leq e^{\varepsilon} \cdot \mathbb{P}\left[\vec{M}\left(x'\right) = \vec{y}\right]$.

Next, we apply union bounds. For any event $E \subseteq \mathcal{Y}$, we decompose the probabilities into that of the corresponding singleton events. (note that the singleton events are mutually exclusive)

$$\mathbb{P}\left[\vec{M}\left(x\right) \in E\right] = \mathbb{P}\left[\vec{M}\left(x\right) \in \bigcup_{\vec{y}_i \in E} \vec{y}_i\right] = \sum_{\vec{y}_i \in E} \mathbb{P}\left[\vec{M}\left(x\right) = \vec{y}_i\right]$$

$$\mathbb{P}\left[\vec{M}\left(x'\right) \in E\right] = \mathbb{P}\left[\vec{M}\left(x'\right) \in \bigcup_{\vec{y}_i \in E} \vec{y}_i\right] = \sum_{\vec{y}_i \in E} \mathbb{P}\left[\vec{M}\left(x'\right) = \vec{y}_i\right]$$

Since we already know $\mathbb{P}\left[\vec{M}\left(x\right) = \vec{y}\right] \leq e^{\varepsilon} \cdot \mathbb{P}\left[\vec{M}\left(x'\right) = \vec{y}\right]$ for all $\vec{y}$, we can add multiple such inequalities together.

$$\sum_{\vec{y}_i \in E} \mathbb{P}\left[\vec{M}\left(x\right) = \vec{y}_i\right] \leq \sum_{\vec{y}_i \in E} e^{\varepsilon} \cdot \mathbb{P}\left[\vec{M}\left(x'\right) = \vec{y}_i\right]$$

$$\sum_{\vec{y}_i \in E} \mathbb{P}\left[\vec{M}\left(x\right) = \vec{y}_i\right] \leq e^{\varepsilon} \sum_{\vec{y}_i \in E} \mathbb{P}\left[\vec{M}\left(x'\right) = \vec{y}_i\right]$$

$$\mathbb{P}\left[\vec{M}\left(x\right) \in E\right] \leq e^{\varepsilon} \cdot \mathbb{P}\left[\vec{M}\left(x'\right) \in E\right]$$

Therefore, $\vec{M}\left(x\right)$ satisfies ε-differential privacy. □

We will now apply this mechanism to `Prio3Histogram` and `Prio3SumVec` in turn.

First, let the `length` parameter of `Prio3Histogram` be denoted by l. Then, each measurement making up a dataset is an element of $\mathcal{X} = \{0, 1, 2, \ldots, l-1\}$, and the query result is a vector of counts, in $\mathcal{Y} = \mathbb{Z}^l_{\geq 0}$. The VDAF's aggregation function is our query function, $f(x)$. It maps each measurement to a one-hot vector, with the position of the one determined by the measurement, and adds them up. The global sensitivity of this query function is $GS_f = 2$. When one measurement in a dataset is replaced with another, then either the result is unchanged, or one count is decreased by one and another is increased by one. Thus, the `scale` parameter is $t = 2/\varepsilon$, and the mechanism will add noise drawn independently from $\text{Lap}_{\mathbb{Z}}(2/\varepsilon)$ to each counter in both aggregate shares.

Let the `length` parameter of `Prio3SumVec` be denoted by l, and the `bits` parameter be denoted by b. Each measurement making up a dataset is an element of $\mathcal{X} = \left\{0, 1, 2, \ldots, 2^b - 1\right\}^l$. The query result is a vector of sums, in $\mathcal{Y} = \mathbb{Z}^l_{\geq 0}$. The VDAF's aggregation function is our query function, $f(x) = \sum_{i=1}^{n} x_i$. The global sensitivity of this query function is $GS_f = \left(2^b - 1\right) \cdot l$, because substituting one measurement may increase or decrease each component of the vector sum by up to $2^b - 1$. Thus, the `scale` parameter is $t = \frac{\left(2^b-1\right)l}{\varepsilon}$, and the mechanism will add noise drawn independently from $\text{Lap}_{\mathbb{Z}}\left(\frac{\left(2^b-1\right)l}{\varepsilon}\right)$ to each sum in both aggregate shares.

References

[1] Canonne, C. L., Kamath, G., and Steinke, T., "The Discrete Gaussian for Differential Privacy", 2020, <https://arxiv.org/abs/2004.00010>.

[2] Nissim, K., Raskhodnikova, S., and Smith, A., "Smooth sensitivity and sampling in private data analysis", 2007, <https://cs-people.bu.edu/ads22/pubs/NRS07/NRS07-full-draft-v1.pdf>.