Tight Differential Privacy Blanket for Shuffle Model

Sayan Biswas[1,2], Kangsoo Jung[1,2], and Catuscia Palamidessi[1,2]

[1] *INRIA, France*
[2] *LIX, École Polytechnique, France*
sayan.biswas@inria.fr, gangsoo.zeong@inria.fr, catuscia@lix.polytechnique.fr

Abstract

With the recent bloom of focus on digital economy, the importance of personal data has seen a massive surge of late. Keeping pace with this trend, the model of data market is starting to emerge as a process to obtain high-quality personal information in exchange of incentives. To have a formal guarantee to protect the privacy of the sensitive data involved in digital economy, *differential privacy (DP)* is the go-to technique, which has gained a lot of attention by the community recently. However, it is essential to optimize the privacy-utility trade-off by ensuring the highest level of privacy protection is ensured while preserving the utility of the data. In this paper, we theoretically derive sufficient and necessary conditions to have tight (ϵ, δ)-DP blankets for the shuffle model, which, to the best of our knowledge, have not been proven before, and, thus, characterize the best possible DP protection for shuffle models which can be implemented in data markets to ensure privacy-preserving trading of digital economy.

1 Introduction and background

As mentioned in [1], privacy in digital economy is critical, especially for end-users who share their personal data. Differential privacy is a de-facto standard for privacy protection, however, it deteriorates the data utility. This trade-off between privacy and utility is a long standing problem in differential privacy.

An intermediate paradigm between the central and the local models of differential privacy (DP), known as the *shuffle model* [2], has recently gained popularity. As an initial step, the shuffle model uses a local mechanism to perturb the data individually like the local model of DP. After this local sanitization, a shuffler uniformly permutes the noisy data to dissolve their links with the corresponding data providers. This allows the shuffle model to achieve a certain level of DP guarantee using less noise than the local model ensuring that the shuffle model provides better utility than the local model whilst retaining the same advantages. Thus, the shuffle model has an advantage in the trade-off between privacy and utility for the digital economy.

The privacy guarantees provided by the shuffle model have been rigorously studied by community of late and various results have been derived, both *analytical* and *numerical*. Obviously, analytical bounds have the advantage that they provide a concrete basis for reasoning and mathematically analysing properties such as privacy-utility trade-off. However, in the case of the shuffle model, most analytical bounds found in the literature are far from being tight. In this paper, we cover this gap and derive tight necessary and sufficient condition for having the tightest (ϵ, δ)-bounds for the DP guarantee provided by the shuffle model with the *k-randomized response* (*k*-RR) local mechanism. We combine the idea of privacy blankets introduced by Balle et al. in [3] and the concept of (ϵ, δ)-adaptive differential privacy (ADP) proposed by Sommer et al. in [4].

Definition 1.1 (Differential privacy [5, 6])**.** For a certain query, a randomizing mechanism $\mathcal{K}$ taking datasets as input, provides (ϵ, δ)-*differential privacy (DP)* if for all neighbouring [1] datasets, D_1 and D_2, and all $S \subseteq \mathrm{Range}(\mathcal{K})$, we have

$$\mathbb{P}[\mathcal{K}(D_1) \in S] \leq e^\epsilon \, \mathbb{P}[\mathcal{K}(D_2) \in S] + \delta$$

Definition 1.2 (Adaptive differential privacy [4])**.** Let us fix $x_0, x_1 \in \mathcal{X}$, where $\mathcal{X}$ is the alphabet of the original (non-privatized) data, and let us fix a member u in the dataset. For a certain query, a randomizing mechanism $\mathcal{K}$ provides (ϵ, δ)-*adaptive differential privacy (ADP) for x_0 and x_1* if for all datasets, $D(x_0)$ and $D(x_1)$, and all $S \subseteq \mathrm{Range}(\mathcal{K})$, we have

$$\mathbb{P}[\mathcal{K}(D(x_0)) \in S] \leq e^\epsilon \, \mathbb{P}[\mathcal{K}(D(x_1) \in S] + \delta$$

where $D(x_0)$ and $D(x_1)$ are datasets differing only in the entry of the fixed member u: $D(x)$ means that u reports x for every $x \in \mathcal{X}$, keeping the entries of all the other users the same.

Definition 1.3 (Tight DP (or ADP) [4])**.** Let $\mathcal{K}$ be an (ϵ, δ)-DP (or ADP for chosen $x_0, x_1 \in \mathcal{X}$) mechanism. We say that δ is *tight* for $\mathcal{K}$ (wrt ϵ and x_0, x_1 in case of ADP) if there is no $\delta' < \delta$ such that $\mathcal{K}$ is (ϵ, δ')-DP (or ADP for x_0, x_1).

Definition 1.4 (Shuffle model [7, 8])**.** Let $\mathcal{X}$ and $\mathcal{Y}$ be discrete alphabets for the original and the noisy data respectively. For any dataset of size $n \in \mathbb{N}$, the *shuffle model* is defined as $\mathcal{M} : \mathcal{X}^n \mapsto \mathcal{Y}^n$, $\mathcal{M} = \mathcal{S} \circ \mathcal{R}^n$, where

i) $\mathcal{R} : \mathcal{X} \mapsto \mathcal{Y}$ is a local randomizer, stochastically mapping each element of the input dataset, sampled from $\mathcal{X}$, onto an element in $\mathcal{X}$, providing ϵ_0-local differential privacy.

ii) $\mathcal{S} : \mathcal{Y}^n \mapsto \mathcal{Y}^n$ is a shuffler that uniformly permutes the finite set of messages of size $n \in \mathbb{N}$, that it takes as an input.

---

[1] differing in exactly one entry

2 Tight differential privacy blanket analysis

In [3] Balle et al. introduced the concept of *privacy blankets* for shuffle models, deriving some quantifiable privacy estimates for them in terms of (ϵ, δ)-DP. In this work, a notion of *strong adversaries* was assumed, and the authors derived a precise condition for having $\epsilon \in \mathbb{R}^+$ and a corresponding space of $\delta \in \mathbb{R}^+$ that ensure a (ϵ, δ)-DP guarantee (Theorem 5.3 of [3]. Although this work gives a straightforward and sufficient condition for having a space of ϵ and δ to encase the shuffle model with diffenrential privacy guarantees, it is crucial to note that no explicit theoretical condition exists that can ensure the tightness[2] of such privacy blankets.

In the same year, Sommer et al. in [4] gave an explicit condition for having a δ for a chosen ϵ that would foster a tight (ϵ, δ)-ADP, for a couple of given inputs, on any probabilistic mechanism (Lemma 5 in [4]).

For comparing utilities between the shuffle model and other privacy mechanisms under a certain level of privacy, having a slack (ϵ, δ)-privacy coverage would not suffice as we would not be able to guage the best privacy guarantee provided without having the tight bounds. Therefore, in this paper, we have investigated and analysed the explicit conditions under which the result of Balle et al. in [3]) improves and ensures a tight privacy blanket for the shuffle model using k-RR mechanism as its local randomizer. In particular, drawing parallels between Lemma 5 of [4], we derived a necessary and sufficient condition for Theorem 5.3 of [3] to ensure a tight (ϵ, δ)-DP guarantee.

We analyze the two cases of Lemma 5 in [4] separately:

$$\textit{Case 1:} \quad \frac{1}{e^{\epsilon_0}} \leq \frac{(e^\epsilon - 1)^2}{(e^\epsilon + 1)^2 (e^{\epsilon_0} - e^{-\epsilon_0})^2} \tag{1}$$

$$\textit{Case 2:} \quad \frac{1}{e^{\epsilon_0}} > \frac{(e^\epsilon - 1)^2}{(e^\epsilon + 1)^2 (e^{\epsilon_0} - e^{-\epsilon_0})^2} \tag{2}$$

Setting $C = 1 - e^{-2} \approx 0.86$, Lemma 5 in [4]) proves that $\mathcal{M}$ guarantees (ϵ, δ)-DP for every choice of $\epsilon \in \mathbb{R}^+$ and $\delta \in \mathbb{R}^+$ satisfying:

$$\delta \geq \begin{cases} \frac{(e^\epsilon+1)^2(e^{\epsilon_0}-e^{-\epsilon_0})^2}{4n(e^\epsilon-1)} e^{-Cn\frac{1}{e^{\epsilon_0}}} & \text{for } \textit{Case 1.} \\ \frac{(e^\epsilon+1)^2(e^{\epsilon_0}-e^{-\epsilon_0})^2}{4n(e^\epsilon-1)} e^{-Cn\frac{(e^\epsilon-1)^2}{(e^\epsilon+1)^2(e^{\epsilon_0}-e^{-\epsilon_0})^2}} & \text{for } \textit{Case 2.} \end{cases} \tag{3}$$

Note that for a given $\epsilon > 0$, if δ_1 and δ_2 both satisfy Lemma 5 [4], then $\max\{\delta_1, \delta_2\}$ will provide a slack privacy blanket, making it hard to quantify a definite differential privacy level for $\mathcal{M}$. It is imperative to have a tight privacy parameter in order to proceed with any sort of privacy-utility analysis and lay down a notion of comparison between $\mathcal{M}$ and other DP mechanisms. Hence, to investigate the existence of a tight δ for a chosen ϵ and to examine the precise conditions under which Result 1 by Balle et al. can provide a tight differential privacy guarantee for $\mathcal{M}$, we consider the minimum δ satisfying (3) for a given ϵ and equate it to Lemma 5 of [4]) and maximize it over all pairs of input values to translate the ADP guarantee of Sommer et al.'s result to the standard notion of DP.

---

[2]as in Definition 1.3

As we are interested to examine if we can find $\epsilon > 0$ and, correspondingly, $\delta > 0$ that provide a tight DP guarantee for $\mathcal{M}$, we define the constants $\kappa_i, i \in \{1, \ldots, 8\}$, that will become handy to simplify the mathematical results derived and used in the paper, as follows:

$$\kappa_1 := \frac{(e^{\epsilon_0} - e^{-\epsilon_0})^2 e^{Cne^{-\epsilon_0}}}{4} \tag{4}$$

$$\kappa_2 := 1 + \frac{e^{\epsilon_0} - 1}{n + (e^{\epsilon_0} - 1)\pi(x_0)(n-1)} \tag{5}$$

$$\kappa_3 := \frac{n + (e^{\epsilon_0} - 1)(n\pi(x_0) + 1 - \pi(x_0))}{e^{\epsilon_0} + k - 1} \tag{6}$$

$$\kappa_4 := 2\operatorname{artanh}\left(\frac{2\sinh(\epsilon_0)}{e^{\epsilon_0/2}}\right) \tag{7}$$

$$\kappa_5 := \frac{\sinh^2 \epsilon_0}{n} \tag{8}$$

Definition 2.1 (Critical Polynomial). For a given privacy parameter, ϵ_0, of the k-RR mechanism used in $\mathcal{M}$ and n samples drawn from $\mathcal{X}$ following a distribution π, let $P_1(x), P_2(x) \in \mathbb{R}[x]$ be polynomials defined as:

$$P_1(x) := \kappa_1 \kappa_2 (x+1)^2, \; P_2(x) := \kappa_3(\kappa_2 - x)(x - 1)$$

Let $P(x) := P_1(x) - P_2(x)$. We call $P(x)$ to be our *critical polynomial*.

Definition 2.2 (Critical Equation). Let the *critical equation* $H(x)$ be defined as:

$$H(x) := 2\kappa_5 \kappa_2 e^{-\frac{Cx^2}{4\kappa_5}} + x^2 \kappa_3(\kappa_2 + 1) - x\kappa_3(\kappa_2 - 1)$$

Theorem 2.1. For $\epsilon > \max\{\kappa_4, \ln \kappa_2\}$, taking δ as in (3) asymptotically provides tight (ϵ, δ)-ADP guarantee for $\mathcal{M}$ wrt x_0, x_1 as $n \to \infty$.

Theorem 2.2. For $\ln \kappa_2 < \epsilon < \kappa_4$, taking δ as in (3) asymptotically provides a tight (ϵ, δ)- ADP blanket for $\mathcal{M}$ wrt x_0, x_1, as $n \to \infty$.

For any $x_i \in \mathcal{X}$ as the primary input of u, let $\mathcal{S}_1(x_i)$ denote the space $[\kappa_4, \ln \kappa_2(x_i)) \subset \mathbb{R}$ and $\mathcal{S}_2(x_i)$ denote the space $(0, \min\{\kappa_4, \ln \kappa_2(x_i)\}) \subset \mathbb{R}$.

Theorem 2.3.   a) If $\kappa_4 < \ln \kappa_2(x_i)$ for all $x_i \in \mathcal{X}$, and if $P(x) = 0$ has a real solution in $[e^{\kappa_4}, \kappa_2)$, then $\mathcal{S}_1 := \bigcap_{x_i \in \mathcal{X}} \mathcal{S}_1(x_i) \neq \phi$. Moreover, for every $\epsilon \in \mathcal{S}_1$, choosing δ as in (3) ensures a tight (ϵ, δ)-DP blanket for $\mathcal{M}$.

b) If $H(x) = 0$ has a real solution in $(0, \mu)$, where $\mu = \min\left\{\tanh\left(\frac{\kappa_4}{2}\right), \tanh\left(\frac{\ln \kappa_2(x_i)}{2}\right)\right\}$, then setting $\mathcal{S}_2 := \bigcap_{x_i \in \mathcal{X}} \mathcal{S}_2(x_i) \neq \phi$. Moreover, for every $\epsilon \in \mathcal{S}_2$, choosing δ ensures a tight (ϵ, δ)-DP blanket for $\mathcal{M}$.

3 Discussion and future work

The theoretical conditions derived in this work to make Balle et al.'s bounds for privacy blankets of shuffle models give us an analytical insight to the cases when we could obtain the best privacy guarantee for shuffle models using k-RR local randomizers. Studying the space of the tight DP guarantees

for the shuffle models could be a possible breakthrough in this area, as this could then be used to tune the hyperparamters such as the privacy level of the local randomizer (ϵ_0) and the number of samples (n) for implementing the shuffle model in various areas of privacy-preserving data analysis. Naturally, this would also help to maximize the utility for a given level of privacy, once the best DP guarantee for the shuffle model is illustrated, aiding to resolve the privacy-utility trade-off for shuffle models of DP.

We plan on studying more generalized forms of shuffle models using different local randomizers and comparing their utilities with the central models. Also, wish to extend this work in the context of digital markets by implementing shuffle models and exploiting its DP guarantees in the context of differentially private data market and privacy pricing [9, 10]. In [9], Jung et al. proposed a federated data trading framework in which data providers coalesce to form a federations to increase their bargaining power in data trading. An immediate extension would be to analyse the mechanism of [9] under the environment of shuffle models implement by each federation and study a data pricing mechanism based on the privacy amplification by shuffling.

References

[1] C. Wang, N. Zhang, and C. Wang, "Managing privacy in the digital economy," *Fundamental Research*, vol. 1, no. 5, pp. 543–551, 2021.

[2] A. Bittau, Ú. Erlingsson, P. Maniatis, I. Mironov, A. Raghunathan, D. Lie, M. Rudominer, U. Kode, J. Tinnes, and B. Seefeld, "Prochlo: Strong privacy for analytics in the crowd," in *Proceedings of the 26th Symposium on Operating Systems Principles*, 2017, pp. 441–459.

[3] B. Balle, J. Bell, A. Gascón, and K. Nissim, "The privacy blanket of the shuffle model," in *Annual International Cryptology Conference*. Springer, 2019, pp. 638–667.

[4] D. M. Sommer, S. Meiser, and E. Mohammadi, "Privacy loss classes: The central limit theorem in differential privacy," *Proceedings on privacy enhancing technologies*, vol. 2019, no. 2, pp. 245–269, 2019.

[5] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography*, S. Halevi and T. Rabin, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 265–284.

[6] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Advances in Cryptology - EUROCRYPT 2006*, S. Vaudenay, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 486–503.

[7] A. Cheu, A. Smith, J. Ullman, D. Zeber, and M. Zhilyaev, "Distributed differential privacy via shuffling," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2019, pp. 375–403.

[8] Ú. Erlingsson, V. Feldman, I. Mironov, A. Raghunathan, K. Talwar, and A. Thakurta, "Amplification by shuffling: From local to central differential privacy via anonymity," in *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2019, pp. 2468–2479.

[9] K. Jung, S. Biswas, and C. Palamidessi, "Establishing the price of privacy in federated data trading," in *Protocols, Strands, and Logic*. Springer, 2021, pp. 232–250.

[10] S. Biswas, K. Jung, and C. Palamidessi, "An incentive mechanism for trading personal data in data markets," in *International Colloquium on Theoretical Aspects of Computing*. Springer, 2021, pp. 197–213.