Enhanced Privacy Bound for Shuffle Model with Personalized Privacy

Yixuan Liu*
Renmin University of China
Beijing, China
liuyixuan@ruc.edu.cn

Yuhan Liu
Renmin University of China
Beijing, China
liuyh2019@ruc.edu.cn

Li Xiong
Emory University
Atlanta, USA
lxiong@emory.edu

Yujie Gu
Kyushu University
Fukuoka, Japan
gu@inf.kyushu-u.ac.jp

Hong Chen
Renmin University of China
Beijing, China
chong@ruc.edu.cn

Abstract

The shuffle model of Differential Privacy (DP) is an enhanced privacy protocol which significantly amplifies the central DP guarantee by anonymizing and shuffling the local randomized data. Yet, deriving a tight privacy bound is challenging due to its complicated randomization protocol. While most existing works focused on uniform local privacy settings, this work focuses on a more practical personalized privacy setting. To bound the privacy after shuffling, we need to capture the probability of each user generating clones of the neighboring data points and quantify the indistinguishability between two distributions of the number of clones on neighboring datasets. Existing works either inaccurately capture the probability or underestimate the indistinguishability. We develop a more precise analysis, which yields a general and tighter bound for arbitrary DP mechanisms. Firstly, we derive the clone-generating probability by hypothesis testing, which leads to a more accurate characterization of the probability. Secondly, we analyze the indistinguishability in the context of f-DP, where the convexity of the distributions is leveraged to achieve a tighter privacy bound. Theoretical and numerical results demonstrate that our bound remarkably outperforms the existing results in the literature. The code is publicly available at https://github.com/Emory-AIMS/HPS.git.

CCS Concepts

• **Security and privacy** → **Privacy protections**.

Keywords

Differential Privacy, Shuffle Model, Personalized Privacy

ACM Reference Format:

Yixuan Liu, Yuhan Liu, Li Xiong, Yujie Gu, and Hong Chen. 2024. Enhanced Privacy Bound for Shuffle Model with Personalized Privacy. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24), October 21–25, 2024, Boise, ID, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3627673.3679911

*Work done while visiting Emory University

CIKM '24, October 21–25, 2024, Boise, ID, USA.

ACM ISBN 979-8-4007-0436-9/24/10
https://doi.org/10.1145/3627673.3679911

1 Introduction

The shuffle model [2] with Differential Privacy (DP) [7] is an advanced privacy protection protocol for distributed data analysis [4, 13, 21, 25]. An intermediate trusted server *shuffler* is introduced between local randomizer [10] and central analyzer [8]. By permuting locally randomized data before sending to the central analyzer, the shuffler brings extra randomness with a privacy amplification effect, i.e., central privacy guarantee after shuffling is significantly stronger than the original local privacy achieved by perturbation.

Many efforts have been put on converting the randomness to a formal privacy guarantee [1, 9, 11, 12, 14]. While most studies achieve privacy bound by assuming a uniform privacy level for all users, this work focuses on a more practical but less studied setting with personalized privacy, where users have different privacy levels on local data points due to different policies or privacy preferences [17–19, 24]. Fig. 1 shows the personalized setting where local data point x_i is associated with a personalized privacy level ϵ_i, δ_i.

A classic privacy analysis for shuffle model amplifies the privacy by leveraging the confounding effect of clones of neighboring data points generated by each user [11, 12]. Specifically, for any neighboring datasets that differ by x_1, the noisy data point from each user could generate a clone of randomized x_1 with a certain probability p. The clones together help to hide the existence of x_1; then the difference of the number of clones on neighboring datasets is estimated for final privacy bounds.

However, driving the probability p and the difference of number of clones is challenging, especially under Personalized Local Differential Privacy (PLDP). Approximating p with the conventional way that reduces any DP randomized mechanism to the worst-case randomized response leads to inaccurate results. Additionally, various privacy parameters exaggerate the complexity of the overall distributions of the number of clones. Existing works [3, 24] approximating it by central limit theorem cause relaxations on privacy bound, especially when the number of users is not large enough.

Motivated by this, we develop a more precise analysis on privacy amplification of shuffle model under both pure- and approximate-PLDP for arbitrary local randomizers. Firstly, we quantify different p contributed by each user with personalized privacy parameters in

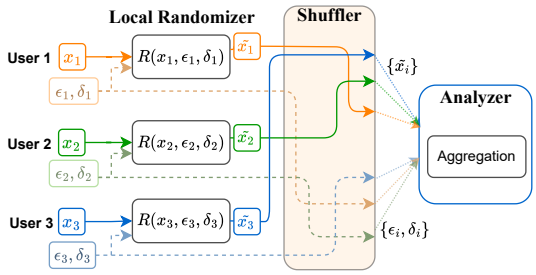

Figure 1: Procedure of shuffle model with personalized privacy. Each user data x_i is randomized locally. Privacy parameters (ϵ_i, δ_i) and perturbed $\tilde{x}_i$ are shuffled. Analyzer aggregates $\tilde{x}_i$ for further statistics or model training.

a more accurate manner. In specific, p is derived by conducting hypothesis testing on the distribution of current noisy data point and the distribution of noisy neighboring data point. By computing the hypothesis testing error, we accurately identify the probability of any data points being wrongly-recognized as x_1. Our method allows computing p on heterogeneous privacy parameters and arbitrary DP local randomizer. Secondly, we analyze the indistinguishability between two overall distributions of the number of clones in the context of f-DP [6]. We depict the overall distributions by Multi-Bernoulli and Binomial distribution. Then inspired by [26], the convexity of the distribution is further exploited to closely characterize the properties of the overall distributions, thus leading to a tighter upper bound on the privacy after shuffling.

Our main contributions are summarized as follows:

- We provide a more precise analysis for privacy amplification effect on the shuffle model for personalized privacy. Confounding effect of individuals and overall distributions are characterized by analytical expressions, which leads to a tighter privacy bound.
- Our work offers a general method to quantify confounding effect of PLDP with hypothesis testing, which enables our analysis to address arbitrary locally differentially private mechanisms and heterogeneous privacy parameters.
- We verify the proposed analysis with numerical results, which demonstrates that our privacy bound significantly exceeds the SOTAs on both pure- and approximate-PLDP.

2 Preliminaries

2.1 Central and Local Differential Privacy

Differential privacy (DP) [8] provides a rigorous privacy guarantee for raw data by introducing random noises to the computation process. The notion is typically applied in a central setting where a trusted server can access the raw data. For the settings without trusted server, local differential privacy (LDP) [10] is proposed. LDP is capable of providing a stronger privacy guarantee than DP, as it protects data against stronger adversaries who have access to every (perturbed) data point in the dataset. Therefore, it is suitable for distributed data collection or publishing [5, 22, 23, 27]. Yet, LDP also suffers from a dissatisfying data utility due to a large amount of noise injection.

Definition 1 (Differential Privacy). *For any $\epsilon, \delta \geq 0$, a randomized algorithm $R : \mathcal{D} \to \mathcal{Z}$ is (ϵ, δ)-DP if for any neighboring datasets $D, D' \in \mathcal{D}$ and any subsets $S \subseteq \mathcal{Z}$, $\Pr[R(D) \in S] \leq e^{\epsilon} \Pr[R(D') \in S] + \delta$.*

Definition 2 (Local Differential Privacy). *For any $\epsilon, \delta \geq 0$, a randomizer $R : \mathcal{D} \to \mathcal{Z}$ is (ϵ, δ)-LDP if $\forall x, x' \in \mathcal{D}$ and $\forall z \in \mathcal{Z}$, $\Pr[R(x) = z] \leq e^{\epsilon} \Pr[R(x') = z] + \delta$.*

ϵ denotes the privacy level, the lower the stronger privacy. δ denotes the failure probability of the randomizer. $\delta = 0$ is pure-LDP, and $\delta > 0$ is approximate-LDP.

2.2 Shuffle-based Privacy

Shuffle model [2] is proposed to strengthen central privacy while preserving local user privacy. Given dataset D, each $x_i \in D$ owned by user i is perturbed locally by a randomizer R to ensure $(\epsilon_i^l, \delta_i^l)$-LDP and sent to shuffler. Shuffler S, a trusted third party, permutes all data points and sends them to an untrusted analyzer A for further computation. Based on the anonymity from shuffling, existing works obtain a strong privacy amplification effect. Most works [1, 9, 11, 12, 14, 20, 26] focus on uniform local privacy setting, Feldman et al. [12] improves privacy bound by generating clones from neighbor data points. Wang et al. [26] applies f-DP and achieves a tighter bound under uniform LDP. As a more common and practical setting, some works [3, 24] focus on personalized settings, while leaving a loose privacy bound due to reduction or approximation.

3 Privacy Analysis

In this section, we first introduce the confounding effect, which captures the randomness introduced by shuffler and serves as the foundation of amplification effect analysis. Then we provide an analytical expression of confounding effect with hypothesis testing, which yields a precise description and results in a stronger amplification effect. At last, we develop our analysis in the context of f-DP. By exploiting the convexity of the mixed distribution generated by the shuffler, we further derive a tighter bound.

3.1 Confounding Effect p

We consider neighboring data points x_1^0 and $x_1^1 \in D$. As noted in [11], after perturbing and shuffling each data point, the output of randomizer on each data point could be seen as samples from the output distribution of randomizer on x_1^0 or x_1^1 with certain probability. And each local randomizer $R(x, \epsilon) : D \to Z$ can also be represented as: $R(x_1^0) = (1-p)Q(x_1^0) + pQ(x_1^1)$ and $R(x_1^1) = pQ(x_1^0) + (1-p)Q(x_1^0)$, where $Q : \{x_1^0, x_1^1\} \to S$ is a randomized algorithm. Hence the following decomposition is given by [11]:

$$R(x_1^0) = e^{\epsilon} pQ(x_1^0) + pQ(x_1^1), \quad R(x_1^1) = pQ(x_1^0) + e^{\epsilon} pQ(x_1^1) \quad (1)$$

$$\forall i \in [2, n], \quad R(x_i) = pR(x_1^0) + pR(x_1^1) + (1-2p)LO_i \quad (2)$$

where LO_i is the leftover distributions. The decomposition above suggests that each output from $R(x_i)$ could be wrongly recognized as coming from x_1^0 or x_1^1 with probability p. In other word, p is the *confounding effect* of $R(x_i)$ on x_1^b, where $b = 0$ or 1, and stronger privacy is achieved with a larger p. Existing works derive p by reducing the LDP mechanism to random response [15], which underestimates the confounding effect of most LDP mechanisms.

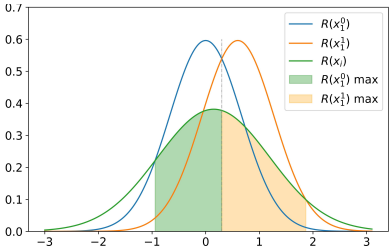

Figure 2: Green area represents $\Pr[R(x_i) \in U_0]$, output of $R(x_i)$ is wrongly recognized from x_1^0; Yellow area represents $\Pr[R(x_i) \in U_1]$, output of $R(x_i)$ is mistaken from x_1^1.

In this work, we achieve a precise p. By conducting hypothesis testing on distributions $R(x_i)$ and $R(x_1^b)$ [1], the type I error captures the probability of wrongly recognizing the output of $R(x_i)$ as an output of $R(x_1^b)$, which is exactly p. For clarity, we derive the value of p for the neighboring data point x_1^b and the rest data point x_i in Section 3.2.1 and 3.2.2 respectively.

3.2 Quantifying p with Hypothesis Testing

3.2.1 Hypothesis Testing on Neighboring Data Point x_1^b. In this section, we demonstrate our hypothesis testing based approach for deriving p at x_1^b, where the confounding of $R(x_1^b)$ only depends on the the privacy budget (ϵ_1, δ_1). Given a random output Z from $R(x_1^b)$, we set the hypothesis testing as follows:

$$H_0\text{: } Z \text{ came from } x_1^0, \qquad H_1\text{: } Z \text{ came from } x_1^1.$$

Then we conduct likelihood ratio test by examining the ratio between probability $p_1^0 = \Pr[R(x_1^0) = Z]$ and $p_1^1 = \Pr[R(x_1^1) = Z]$, and reject H_0 when $p_1^0/p_1^1 < 1$. The rejection region is defined as

$$S = \{z | \Pr[R(x_1^0) = z] < \Pr[R(x_1^1) = z]\}.$$

According to Neyman–Pearson lemma [16], likelihood ratio test is the most powerful way to distinguish two distributions. Hence with such S, we achieve the lower bound of p. As for approximate-DP, the privacy protection fails when outputs $z \in T_\delta$ where

$$T_\delta = \{z | \Pr[R(x) = z] < -\delta/2 \text{ or } \Pr[R(x) = z] > 1 - \delta/2\}.$$

After removing the failure set T_δ, the p is lower bounded by

$$\Pr[R(x_1^0) \in S \backslash T_{\delta_1}^0] = \Pr[R(x_1^1) \in \bar{S} \backslash T_{\delta_1}^1] = p_1.$$

where $\bar{S}$ is the complement of S, $T_{\delta_1}^0$ denotes the failure set on x_1^0 with δ_1, $T_{\delta_1}^1$ is on x_1^1 with δ_1. Then Eq.(1) is rewritten as Eq. (3). By further considering the distribution of concrete DP mechanisms, we are able to achieve the exact expression of p.

$$R(x_1^0) = (1-p_1)R(x_1^0) + p_1 R(x_1^1), \quad R(x_1^1) = p_1 R(x_1^0) + (1-p_1)R(x_1^1) \quad (3)$$

3.2.2 Hypothesis Testing on Rest Data Points x_i. We then extend the method to x_i for $i \in [2, n]$. The main difference lies in the confounding effect that involves heterogeneous privacy parameters (ϵ_1, δ_1) and (ϵ_i, δ_i) now. Given a random output Z of $R(x_i, \epsilon_i, \delta_i)$ and $R(x_1^b, \epsilon_1, \delta_1)$, we set hypothesis testing:

$$H_0\text{: } Z \text{ came from } x_i, \qquad H_1\text{: } Z \text{ came from } x_1^b.$$

[1] For convenience, we use the simplified notation $R(x_i)$ instead of $R(x_i, \epsilon_i, \delta_i)$ when it is clear from the context, as (ϵ_i, δ_i) is always binding with x_i.

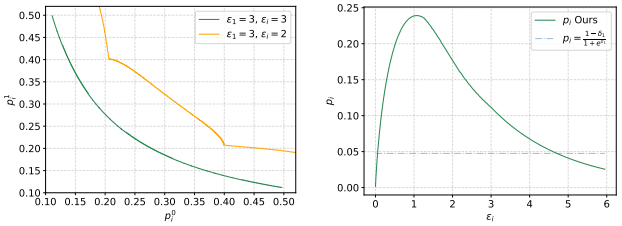

(a) Examples of p_i^0 v.s. p_i^1, x_i varies within $[x_1^0, x_1^1]$. The fluctuation on orange line suggests the different confounding patterns under personalized privacy.

(b) Fix $\epsilon_1 = 3$, p_i changes as ϵ_i varies. p_i is highest at around $\epsilon_1 = 1$, as the testing error maximizes when $R(x_i)$ is almost covered by $R(x_1^0)$ and $R(x_1^1)$.

Figure 3: Confounding effect p under personalized privacy.

Noticing that H_1 indicates Z came from x_1^0 or x_1^1, with likelihood ratio test, we set the rejection region as

$$U = \{z | \Pr[R(x_i) = z] < \max(\Pr[R(x_1^0) = z], \Pr[R(x_1^1) = z])\}$$

Therefore, with $\Pr[R(x_i) \in U]$ null hypothesis is true but rejected, i.e., $R(x_i)$ is wrongly recognized as $R(x_1^b)$. U could be further partitioned into two subsets U_0 and U_1:

$$U_0 = \{z | \Pr[R(x_i) = z] < \Pr[R(x_1^0) = z] \text{ and } \Pr[R(x_1^1) = z] < \Pr[R(x_1^0) = z]\}$$
$$U_1 = \{z | \Pr[R(x_i) = z] < \Pr[R(x_1^1) = z] \text{ and } \Pr[R(x_1^0) = z] < \Pr[R(x_1^1) = z]\}$$

where $U_0 \cup U_1 = U$ (Cf. Fig. 2). Similar with Section 3.2.1, the failure set due to δ_i is removed from U. Accordingly, the probabilities of type I error on x_1^0 and x_1^1 are defined as:

$$p_i^0 = \Pr[R(x_i) \in U_0 \backslash T_{\delta_i}^i], \qquad p_i^1 = \Pr[R(x_i) \in U_1 \backslash T_{\delta_i}^i]$$

where $T_{\delta_i}^i$ represents the failure set on x_i with δ_i.

We observe that p_i^0 and p_i^1 change as x_i changes (Cf. Fig. 3(a)). As $x_1^0 \le x_i \le x_1^1$, the worst-case happens when $x_i = x_1^0$ (or $x_i = x_1^1$). Considering the fact that privacy is breached at the weakest spot, we adopt minimal p_i to describe the confounding effect of $R(x_i)$. Hence Eq.(2) is rewritten with $p_i = \min(p_i^0, p_i^1)$:

$$\forall i \in [2, n], R(x_i) = p_i R(x_1^0) + p_i R(x_1^1) + (1 - 2p_i) LO(x_i). \quad (4)$$

3.3 Privacy Amplification with f-DP

In this section, we achieve a tighter privacy bound of shuffle model under (ϵ_i, δ_i)-PLDP with f-DP.

After deriving p, the clones of $R(x_1^b)$ by shuffling are generated. Based on [11], the overall distributions of number of clones on D_0 and D_1 are denoted as P and Q, with $w = p_1$,

$$P = (1 - w)P_0 + wQ_0 \quad \text{and} \quad Q = (1 - w)Q_0 + wP_0.$$

where $P_0 \sim (A+1, C-A)$, $Q_0 \sim (A, C-A+1)$ with $A \sim Bin(C, 1/2)$. Considering p_i varies under PLDP, we have $C_i \sim Bern(2p_i)$, $C = \sum_{i=1}^{n-1} C_i$. Following the idea in [26] that mixed distributions are more indistinguishable when indices are unknown, the lower bound of trade-off function of overall distribution could be derived by establishing trade-off function on sub-distributions for each possible situation with certain weights. Specifically, P_0 is the mixture of $\{(A_i + 1, i - A_i)\}_{i=0}^{n-1}$ with weights $w_i^0 = \Pr[C = i]$, Q_0 is the mixture of $\{(A_i, i - A_i + 1)\}_{i=0}^{n-1}$ with the same w_i^0 and $A_i \sim Bin(i, 1/2)$. Let f_i, F_i be the probability mass function and distribution function

Table 1: PLDP privacy parameters ϵ^l, δ^l ($\delta^l = 0$ for pure PLDP). $\mathcal{U}, \mathcal{N}$ represent Uniform and Gaussian Distribution respectively. δ^s after shuffling is 10^{-5}.

Name	$\epsilon^l = \{\epsilon_i^l\}_{i\in[n]}$	$\delta^l = \{\delta_i^l\}_{i\in[n]}$	clip range
Uniform1	$\mathcal{U}(0.05, 1)$	$0, 10^{-10}$	$[0.05, 1]$
Gauss1	$\mathcal{N}(0.8, 0.5)$	$0, 10^{-10}$	$[0.05, 1]$
Uniform2	$\mathcal{U}(0.5, 2)$	$0, 10^{-10}$	$[0.5, 2]$
Gauss2	$\mathcal{N}(1.5, 0.5)$	$0, 10^{-10}$	$[0.5, 2]$

of $Bin(i, 1/2)$ respectively. By Lemma 3.1 in [26], we achieve trade-off function f_s under both pure-PLDP (let $\delta_i = 0$) and approximate-PLDP settings ($\delta_i > 0$).

Theorem 1 (Trade-off function). *The trade-off function of shuffling process is defined as $f_s(\alpha(t))$, for $t \geq 0$, each $\alpha(t) = \sum_{i=0}^{n-1} w_i^0 F_i(i - \frac{i+1}{t+1}) \in [0, 1]$. The function f_s at $\alpha(t)$ is*

$$f_s(\alpha(t)) = (1-\delta_1)(2w(1-\alpha(t)) + (1-2w)\sum_{i=0}^{n-1} w_i^0 F_i) + \delta_1(1-\alpha(t))$$

where F_i is the abbreviation of $F_i(i + 1 - \frac{i+1}{t+1})$.

Then we convert it to DP based on primal-dual perspective [6].

Theorem 2 (Enhanced Privacy Bound). *The shuffling process (with randomizer, shuffler, and analyzer) $R \circ S \circ A$ is $(\epsilon, \delta_s(\epsilon)) - DP$ for any $\epsilon > 0$ with*

$$\delta_s(\epsilon) = (-e^\epsilon + (1-\delta_1)2w + \delta_1)[\sum_{i=1}^{n-1} w_i^0 F_i(i - \frac{i+1}{t_\epsilon + 1})] + (1-\delta_1)(1-2w)[\sum_{i=1}^{n-1} w_i^0 F_i(i+1 - \frac{i+1}{t_\epsilon + 1})] \quad (5)$$

where $t_\epsilon = \inf\{t : (1-\delta_1)(-2w + (1-2w)l(t)) - \delta_1 \geq -e^\epsilon\}$, $w = p_1$, $l(t) = -\sum_{i=1}^{n-1} w_1^0 f_i(\lfloor i + 1 - \frac{i+1}{t+1} \rfloor) / \sum_{i=1}^{n-1} w_1^0 f_i(\lfloor i - \frac{i+1}{t+1} \rfloor)$.

Here we bound the worst case: user 1 with x_1^b adopts weakest privacy budget, $\epsilon_1 = \max(\epsilon_i)$. (Considering δ_i is negligible in usual setting, δ_1 is the corresponding parameter).

4 Experiment Results

We show the privacy bound with various personalized privacy settings, and the different number of users.

Experiment Setting. We evaluate several PLDP parameter settings as Tab.1. Baselines include: for pure DP, BBGN [1], FV [12], CCC [3], LZX [24]; for approximate DP, FV [12], CCC [3]. Notice that BBGN and FV lack the analysis on personalized privacy, only the approximate bound is demonstrated by using $\max(\epsilon_i)$ for all data points. We set the same δ^l for all users for convenience, as FV is easy to be unbound with large δ^l. For our bound, we select Laplace Mechanism and Gaussian Mechanism for evaluating pure and approximate-PLDP respectively. In practical applications, our analysis allows any personalized δ_i and local randomizers.

Privacy Amplification with fixed δ^s. Fig.4 provides the numerical evaluations for privacy amplification effect with various PLDP settings and the number of users. We made two observations. (1) Our bound achieves the strongest privacy amplification effect. The results come from a precise p with hypothesis testing on the concrete mechanism, and sharp bound with f-DP. (2) Compared to pure-PLDP, the bound on approximate-PLDP is tighter. It is reasonable from two aspects: first, Gaussian Mechanism is much more noisy (larger variance) than Laplace Mechanism under the same ϵ. Hence the confounding effect p is larger on approximate-PLDP; second, f-DP precisely characterizes the Gaussian distribution, hence the bound is tighter on approximate-PLDP.

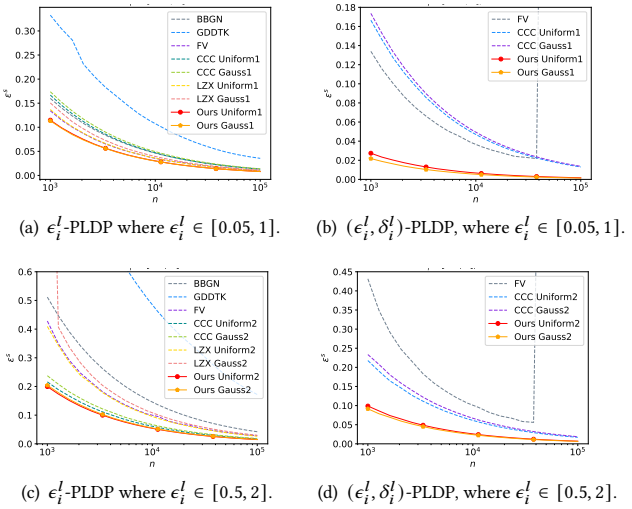

(a) ϵ_i^l-PLDP where $\epsilon_i^l \in [0.05, 1]$. (b) $(\epsilon_i^l, \delta_i^l)$-PLDP, where $\epsilon_i^l \in [0.05, 1]$.

(c) ϵ_i^l-PLDP where $\epsilon_i^l \in [0.5, 2]$. (d) $(\epsilon_i^l, \delta_i^l)$-PLDP, where $\epsilon_i^l \in [0.5, 2]$.

Figure 4: Privacy bounds with various number of data points and privacy parameters, for pure- and approximate-PLDP.

Table 2: δ^s after shuffling comparison under (ϵ^l, δ^l)-PLDP with Uniform2, $n = 10000$.

ϵ^s	0.01	0.03	0.05	0.08	0.1
δ^s [12]	0.0083	0.0029	0.0008	6×10^{-5}	1×10^{-5}
δ^s [3]	0.0042	0.0007	5×10^{-5}	3×10^{-7}	3×10^{-9}
δ^s (Ours)	0.0007	2×10^{-6}	3×10^{-10}	1×10^{-18}	1×10^{-25}

Privacy Amplification with fixed ϵ^s. Tab.2 presents values of δ^s after shuffling with different fixed ϵ^s values. Due to limited space we only show the result of approximate-PLDP, the performance on pure-PLDP is similar. Notably, under the same ϵ^s, our bound on δ is significantly smaller than baselines.

5 Conclusion

This work achieves a refined privacy bound on shuffle model for both pure- and approximate-PLDP. To tighten the bound, we provide a full analysis on the confounding effect of perturbed individual data point and the overall distributions. Our bound on ϵ is up to 5 times tighter than SOTAs.

Acknowledgments

This research has been funded in part by the National Key Research & Develop Plan 2023YFB4503600; National Natural Science Foundation of China (NSFC) U23A20299, 62072460, 62172424, 62276270, 62322214; National Science Foundation (NSF) CNS-2124104, CNS-2125530, IIS-2302968; National Institute of Health (NIH) R01LM013712, R01ES033241. The corresponding author is Hong Chen.

References

[1] Borja Balle, James Bell, Adrià Gascón, and Kobbi Nissim. 2019. The privacy blanket of the shuffle model. In *Annual International Cryptology Conference*. Springer, 638–667.

[2] Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. 2017. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th symposium on operating systems principles*. 441–459.

[3] E Chen, Yang Cao, and Yifei Ge. 2024. A Generalized Shuffle Framework for Privacy Amplification: Strengthening Privacy Guarantees and Enhancing Utility. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 11267–11275.

[4] Wei-Ning Chen, Dan Song, Ayfer Ozgur, and Peter Kairouz. 2024. Privacy amplification via compression: Achieving the optimal privacy-accuracy-communication trade-off in distributed mean estimation. *Advances in Neural Information Processing Systems* 36 (2024).

[5] Graham Cormode, Somesh Jha, Tejas Kulkarni, Ninghui Li, Divesh Srivastava, and Tianhao Wang. 2018. Privacy at scale: Local differential privacy in practice. In *Proceedings of the 2018 International Conference on Management of Data*. 1655–1658.

[6] Jinshuo Dong, Aaron Roth, and Weijie J Su. 2022. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 84, 1 (2022), 3–37.

[7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*. Springer, 265–284.

[8] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (2014), 211–407.

[9] Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. 2019. Amplification by shuffling: From local to central differential privacy via anonymity. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2468–2479.

[10] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. 1054–1067.

[11] Vitaly Feldman, Audra McMillan, and Kunal Talwar. 2022. Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by shuffling. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 954–964.

[12] Vitaly Feldman, Audra McMillan, and Kunal Talwar. 2023. Stronger privacy amplification by shuffling for rényi and approximate differential privacy. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 4966–4981.

[13] Antonious Girgis, Deepesh Data, Suhas Diggavi, Peter Kairouz, and Ananda Theertha Suresh. 2021. Shuffled model of differential privacy in federated learning. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2521–2529.

[14] Antonious M Girgis and Suhas Diggavi. 2024. Multi-message shuffled privacy in federated learning. *IEEE Journal on Selected Areas in Information Theory* 5 (2024), 12–27.

[15] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. 2015. The composition theorem for differential privacy. In *International conference on machine learning*. PMLR, 1376–1385.

[16] Erich L Lehmann and George Casella. 2006. *Theory of point estimation*. Springer Science & Business Media.

[17] Junxu Liu, Jian Lou, Li Xiong, Jinfei Liu, and Xiaofeng Meng. 2021. Projected federated averaging with heterogeneous differential privacy. *Proceedings of the VLDB Endowment* 15, 4 (2021), 828–840.

[18] Junxu Liu, Jian Lou, Li Xiong, Jinfei Liu, and Xiaofeng Meng. 2024. Cross-silo Federated Learning with Record-level Personalized Differential Privacy. *arXiv preprint arXiv:2401.16251* (2024).

[19] Junxu Liu, Jian Lou, Li Xiong, and Xiaofeng Meng. 2023. Personalized Differentially Private Federated Learning without Exposing Privacy Budgets. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 4140–4144.

[20] Ruixuan Liu, Yang Cao, Hong Chen, Ruoyang Guo, and Masatoshi Yoshikawa. 2021. Flame: Differentially private federated learning in the shuffle model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 8688–8696.

[21] Yixuan Liu, Hong Chen, Yuhan Liu, and Cuiping Li. 2021. Privacy-preserving techniques in federated learning. *Journal of Software* 33, 3 (2021), 1057–1092.

[22] Yuhan Liu, Tianhao Wang, Yixuan Liu, Hong Chen, and Cuiping Li. 2024. Edge-Protected Triangle Count Estimation under Relationship Local Differential Privacy. *IEEE Transactions on Knowledge and Data Engineering* (2024).

[23] Yuhan Liu, Suyun Zhao, Yixuan Liu, Dan Zhao, Hong Chen, and Cuiping Li. 2022. Collecting triangle counts with edge relationship local differential privacy. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2008–2020.

[24] Yixuan Liu, Suyun Zhao, Li Xiong, Yuhan Liu, and Hong Chen. 2023. Echo of neighbors: privacy amplification for personalized private federated learning with shuffle model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 11865–11872.

[25] Mary Scott, Graham Cormode, and Carsten Maple. 2022. Aggregation and transformation of vector-valued messages in the shuffle model of differential privacy. *IEEE Transactions on Information Forensics and Security* 17 (2022), 612–627.

[26] Chendi Wang, Buxin Su, Jiayuan Ye, Reza Shokri, and Weijie Su. 2024. Unified Enhancement of Privacy Bounds for Mixture Mechanisms via f-Differential Privacy. *Advances in Neural Information Processing Systems* 36 (2024).

[27] Ning Wang, Xiaokui Xiao, Yin Yang, Jun Zhao, Siu Cheung Hui, Hyejin Shin, Junbum Shin, and Ge Yu. 2019. Collecting and analyzing multidimensional data with local differential privacy. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 638–649.