



HAL
open science

Preventing Manipulation Attack in Local Differential Privacy Using Verifiable Randomization Mechanism

Fumiyuki Kato, Yang Cao, Masatoshi Yoshikawa

► **To cite this version:**

Fumiyuki Kato, Yang Cao, Masatoshi Yoshikawa. Preventing Manipulation Attack in Local Differential Privacy Using Verifiable Randomization Mechanism. 35th IFIP Annual Conference on Data and Applications Security and Privacy (DBSec), Jul 2021, Calgary, AB, Canada. pp.43-60, 10.1007/978-3-030-81242-3_3. hal-03677038

HAL Id: hal-03677038

<https://inria.hal.science/hal-03677038>

Submitted on 24 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



This document is the original author manuscript of a paper submitted to an IFIP conference proceedings or other IFIP publication by Springer Nature. As such, there may be some differences in the official published version of the paper. Such differences, if any, are usually due to reformatting during preparation for publication or minor corrections made by the author(s) during final proofreading of the publication manuscript.

Preventing Manipulation Attack in Local Differential Privacy using Verifiable Randomization Mechanism

Fumiyuki Kato¹, Yang Cao², and Masatoshi Yoshikawa³

¹ Graduate School of Informatics, Kyoto University Dept. of Social Informatics,
Japan

fumiyuki@db.soc.i.kyoto-u.ac.jp

² yang@i.kyoto-u.ac.jp

³ yoshikawa.masatoshi.5w@kyoto-u.ac.jp

Abstract. Local differential privacy (LDP) has been received increasing attention as a formal privacy definition without a trusted server. In a typical LDP protocol, the clients perturb their data locally with a randomized mechanism before sending it to the server for analysis. Many studies in the literature of LDP implicitly assume that the clients honestly follow the protocol; however, two recent studies show that LDP is generally vulnerable under malicious clients. Cao et al. (USENIX Security '21) and Cheu et al. (IEEE S&P '21) demonstrated that the malicious clients could effectively skew the analysis (such as frequency estimation) by sending fake data to the server, which is called data poisoning attack or manipulation attack against LDP. In this paper, we propose secure and efficient verifiable LDP protocols to prevent manipulation attacks. Specifically, we leverage Cryptographic Randomized Response Technique (CRRT) as a building block to convert existing LDP mechanisms into a verifiable version. In this way, the server can verify the completeness of executing an agreed randomization mechanism on the client side without sacrificing local privacy. Our proposed method can completely protect the LDP protocol from output manipulation attacks, and significantly mitigates unexpected damage from malicious clients with acceptable computational overhead.

Keywords: Local Differential Privacy, Manipulation Attack, Data Poisoning, Verifiable computation, Oblivious Transfer

1 Introduction

Today's data science has been very successful in collecting and utilizing large amounts of data. Useful data often includes personal information, and there are serious privacy concerns. In particular, recent data breaches [14] [15] and strict rules [12] [13] by the government significantly promote the concerns. Local differential privacy (LDP) [6] [5] is a promising privacy-enhanced technique for collecting sensitive information. Each client perturbs sensitive data locally by a

randomized mechanism satisfying differential privacy. A server can run analysis such as frequency estimation based on the perturbed data without accessing the raw data. We can see the effectiveness and feasibility of LDP in recent production releases of the platformers such as Google [7], Apple [8], and Microsoft [9], which all utilize LDP for privacy-preserved data curation.

While many studies have been focusing on improving LDP’s utility [10] [11] [16] [17] [18] [19] in the literature, recent studies [1] [2] report a vulnerability of LDP protocol and alert the lack of security. Specifically, [1] [2] show that the malicious clients can manipulate the analysis, such as frequency estimation, by sending false data to the server. Malicious clients can skew the estimations effectively by considering that estimations are calculated by normalizing with randomization probability defined in the LDP protocol and can even control the estimations. Their studies significantly highlight the necessity of a secure LDP protocol to defend against malicious clients. The problematic point of protecting such an attack is that, in a general LDP protocol, others cannot verify the integrity of results without the original data. The randomization would provide data providers plausible deniability for their outputs.

To the best of our knowledge, no effective way in the literature can *completely* prevent manipulation attacks. Although Cao et al. [1] showed some of the countermeasures against malicious clients, their empirical results showed that preventing against *output manipulation attack* is still an open problem. Among their proposed methods, the one normalizing the estimated probability distribution was shown to be to some extent effective for input-manipulation (i.e., the attackers can falsely manipulate the raw input data but honestly execute the local randomized mechanism). However, the proposed countermeasures in [1] are not very effective for output-manipulation attack (i.e., the attackers can arbitrarily change the output of the local randomized mechanism). In addition, their detection-based countermeasures are based on the assumption of specific attack methods and may not be effective against arbitrary output-manipulation attacks. The authors concluded the need for more robust defenses against these attacks. Concurrently, Cheu et al. [2] also emphasize the same conclusion for manipulation attacks they call. There is another promising direction against an attacker who exploits the random mechanism of Differential privacy. Narayan et al. [3] propose an interesting scheme to prove integrity for executing correct randomization mechanisms for Differential privacy. However, their setting is different from ours since they focus on central DP with the data curator, who has the sensitive data, and the analyst, who creates the proof (in our setting, the client needs to prove their local execution).

To solve these problems, we design a novel verifiable LDP protocol based on Multi-Party Computation (MPC) techniques in this work. Our contributions are summarized below. First, we categorize the attacks of malicious clients into two classes, *output-manipulation* and *input-manipulation* (formally defined in Section 3). For input-manipulation attacks, efficient countermeasures have been provided in [1], but existing studies cannot prevent output-manipulation wholly and effectively. We analyze the effectiveness of output-manipulation compared

to input-manipulation, highlight the importance of output-manipulation protection, and formalize the definition of *output-manipulation secure* LDP protocol. Second, we propose secure and efficient verifiable LDP protocols to prevent manipulation attack. Our protocols enable the server to verify the completeness of executing an agreed randomization mechanism on the client side without sacrificing local privacy. Specifically, we leverage Cryptographic Randomized Response Technique (CRR) [4] as a building block to convert existing state-of-the-art LDP mechanisms including kRR [10], OUE [11], and OLH into output-manipulation secure LDP protocols with negligible utility loss. Our proposed secure protocols do not assume any specific attack, and work effectively against general output-manipulation, and thus is more potent than previously proposed countermeasures. Third, we conduct intensive experiments to test the performance of the proposed protocols. We demonstrate that the proposed methods can completely protect the LDP protocol from output manipulation attacks with acceptable computational overhead.

2 Background: Attacks on LDP protocols

2.1 Local Differential Privacy

Differential privacy (DP) [6] is a rigorous mathematical privacy definition, which quantitatively evaluates the degree of privacy protection when we publish outputs about sensitive data in a database. DP is a central model where a trusted server collects sensitive data and releases differentially private statistical information to an untrusted third party. On the other hand, Local DP (LDP) is a local model, considering an untrusted server that collects clients' sensitive data. Clients perturb their data on their local environment and send only randomized data to the server to protect privacy.

In this work, we suppose server S collects data and aggregates them, and N clients c_i ($0 \leq i \leq N - 1$) send their sensitive data in a local differentially private manner. Each client has an item v which is categorical data, and the items have d domains and $v \in [0, d - 1](:= [d])$. Additionally, v_i denotes c_i 's item. The clients randomize their data by randomization mechanism \mathcal{A} , and c_i send $\mathcal{A}(v_i) = y_i (\in D)$ to the server, where D is the output space of \mathcal{A} . The server estimates some statistics by $\mathcal{F}(y_0, \dots, y_{N-1})$. In particular of this work, \mathcal{F}_k corresponds to *frequency estimation* for item k (i.e., how many clients have chosen item k). The formal LDP definition is as follows:

Definition 1 (ϵ -local differential privacy (ϵ -LDP)). *A randomization mechanism \mathcal{A} satisfies ϵ -LDP, if and only if for any pair of input values $v, v' \in [d]$ and for all randomized output $y \in D$, it holds that*

$$\Pr[\mathcal{A}(v) \in y] \leq e^\epsilon \Pr[\mathcal{A}(v') \in y].$$

Under a specific randomized algorithm \mathcal{A} , we want to estimate the frequency of any items. Wang et al. [11] introduce "pure" LDP protocols with nice symmetric property and a generic aggregation procedure to calculate the unbiased

frequency estimations from given randomization probabilities. Let **Support** be a function that maps each possible output y to a set of input that y supports. **Support** is defined for each LDP protocol, and it specifies how the estimation can be computed under the LDP protocol. A formal definition of pure LDP is as follows:

Definition 2 (Pure LDP [11]). *A protocol given by \mathcal{A} and **Support** is pure if and only if there exist two probability values $p > q$ such that for all v_1 ,*

$$\begin{aligned} \Pr[\mathcal{A}(v_1) \in \{o|v_1 \in \text{Support}(o)\}] &= p, \\ \forall_{v_2 \neq v_1} \Pr[\mathcal{A}(v_2) \in \{o|v_1 \in \text{Support}(o)\}] &= q. \end{aligned} \quad (1)$$

where p, q are probabilities, and q must be the same for all pairs of v_1 and v_2 .

While maximizing p and minimizing q make the LDP protocol more accurate, under ϵ -LDP it must be $\frac{p}{q} \leq e^\epsilon$. The important thing is that, in pure LDP protocol, we can simply estimate the frequency of item k as follows:

$$\mathcal{F}_k = \frac{\sum_i \mathbb{1}_{\text{Support}(y^i)}(k) - Nq}{p - q} \quad (2)$$

We can interpret that this formula normalizes observed frequencies using probabilities p and q to adjust for randomization.

For frequency estimation under LDP, we introduce three state-of-the-art randomization mechanisms, kRR [10], OUE [11] and OLH [11]. These mechanisms includes three steps: (1) **Encode** is encoding function: $\mathcal{E} : v(\in [d]) \rightarrow v'(\in [g])$, (2) **Perturbation** is randomized function: $\mathcal{A} : v'(\in [g]) \rightarrow y(\in D)$, (3) **Aggregation** calculates estimations from all collected values: $\mathcal{F} : (y_0, \dots, y_{N-1}) \rightarrow \mathbb{R}$. Formal proofs that each protocol satisfies ϵ -LDP can be found in [11].

k-ary Randomized Response (kRR) is an extension of Randomized Response [32] to meet ϵ -LDP. In particular, kRR provides accurate results in small item domains. This mechanism does not require any special encoding, and provides a identity mapping $\mathcal{E}(v) = v$ ($[g] = [d]$). Perturbation is as follows;

$$\Pr[\mathcal{A}(v) = y] = \begin{cases} p = \frac{e^\epsilon}{e^\epsilon + d - 1}, & \text{if } y = v \\ q = \frac{1 - p}{d - 1} = \frac{1}{e^\epsilon + d - 1}, & \text{if } y \neq v \end{cases} \quad (3)$$

For aggregation, we can consider **Support** function as $\text{Support}(v) = (v)$ and make this follow pure LDP protocol (Def. 2). Therefore, aggregation follows Eq.(2).

Optimized Unary Encoding (OUE) encodes item v into d -length bit vector and encode function is defined as $\mathcal{E}(v) = [0, \dots, 0, 1, 0, \dots, 0]$ where only single bit corresponding to v -th position is 1. Final output space is also d dimensional bit vector $\{0, 1\}^d$ (e.g. $\mathbf{y} = [1, 0, 1, 1, 0]$). Let i -th bit of output vector as y_i ,

perturbation is as follows;

$$Pr[y_i = 1] = \begin{cases} p = \frac{1}{2}, & \text{if } i = v \\ q = \frac{1}{e^\epsilon + 1}, & \text{if } i \neq v \end{cases} \quad (4)$$

These p and q minimize the variance of the estimated frequency in similar bit vector encoding (e.g. RAPPOR [7]). In aggregation step, we consider **Support** function as $Support(\mathbf{y}) = \{v | y_v = 1\}$, and also calculate using Eq.(2).

Optimized Local Hashing (OLH) employs hash function for dimensional reduction to reduce communication costs. It picks up H from a universal hash function family \mathbb{H} , and H maps $v \in [d]$ to $v' \in [g]$ where $2 \leq g < d$. Therefore, encode function is $\mathcal{E}(v) = H(v)$. Perturbation is the same as kRR, except that the input/output space is $[g]$. Then, p and q is defined as follows;

$$Pr[\mathcal{A}(x) = y] = \begin{cases} p = \frac{e^\epsilon}{e^\epsilon + g - 1}, & \text{if } y = H(v) \\ q = \frac{1}{g} \cdot p + \left(1 - \frac{1}{g}\right) \cdot \frac{1}{e^\epsilon + g - 1} = \frac{1}{g}, & \text{if } y \neq H(v) \end{cases} \quad (5)$$

In aggregation step, we consider **Support** function as $Support(\mathbf{y}) = \{v | v \in [d] \text{ and } y = H(v)\}$ and follow Eq.(2) using p and q .

2.2 Attacks on LDP protocols

In this subsection, we introduce two important studies suggesting caution to the necessity of secure LDP protocols.

Targeted Attack. Cao et al. [1] focus on *targeted* attacks to LDP protocols, where the attacker tries to promote the estimated frequencies of a specific item set. Considering the attacker against the LDP protocols, M malicious clients, who can arbitrarily control local environments and send crafted data to the server, are injected by the attacker. (They call *data poisoning attacks*.) Attacker wants to promote r target items $T = \{t_1, \dots, t_r\}$ in the frequency estimation. Cao et al. propose three attacks: Random perturbed-value attack (RPA), Random item attack (RIA), Maximal gain attack (MGA). The first two attacks are as baselines and MGA is an optimized attack. In RIA, malicious clients perform uniform random samplings of a value from the target item set. And then, following the LDP protocol, encoding and perturbation are performed and sent to the server. MGA is more complicated than others. It aims to maximize the attacker's overall gain G : sum of the expected frequency gains for the target items, $G = \sum_{t \in T} \mathbb{E}[\Delta f_t]$ where Δf_t represents the increase of estimated frequency of item t ($\forall t \in T$) from without attack to with attack. In MGA, the output item selection is performed according to the optimal solution maximizing the attacker's gain and sent to the server without perturbation.

	kRR	OUE	OLH
RPA (output-manipulation)	$\beta(\frac{r}{d} - f_T)$	$\beta(r - f_T)$	$-\beta f_T$
RIA (input-manipulation)	$\beta(1 - f_T)$	$\beta(1 - f_T)$	$\beta(1 - f_T)$
MGA (output-manipulation)	$\beta(1 - f_T) + \frac{\beta(d-r)}{e^\epsilon - 1}$	$\beta(2r - f_T) + \frac{2\beta r}{e^\epsilon - 1}$	$\beta(2r - f_T) + \frac{2\beta r}{e^\epsilon - 1}$

Table 1. MGA can achieve the highest **gains** against all three protocols. $\beta = \frac{M}{N+M}$ and $f_T = \sum_{t \in T} f_t$ in the table. (The summary results of [11].)

	kRR	OUE	OLH
Targeted Attack [1]	$+\left(\frac{\beta(d-r)}{e^\epsilon - 1}\right)$	$+\left(\beta(2r-1) + \frac{2\beta r}{e^\epsilon - 1}\right)$	$+\left(\beta(2r-1) + \frac{2\beta r}{e^\epsilon - 1}\right)$
Untargeted Attack [2]	$\times \Omega\left(\frac{\sqrt{d}}{\epsilon}\right)$		

Table 2. Overall, output-manipulations are much more vulnerable than input-manipulation. The differences of both manipulations gain are calculated by output-manipulation gain – input-manipulation gain (resp. output-manipulation gain / input-manipulation gain) in Targeted (resp. Untargeted) Attack.

Cao et al. describe the details of these three attacks against kRR, OUE, OLH in the frequency estimation and give theoretical analysis. The summary of the results is shown in Table 1. The table shows the overall gains of the three attacks against kRR, OUE, and OLH. MGA can achieve the highest gains for all protocols, clearly because MGA maximizes the gains. A notable point is a difference, summarized in Table 2, showing the difference in gains between MGA and RIA. They respectively correspond to *output-manipulation* and *input-manipulation* (described later) in our work. Note that the difference is remarkable, especially under the higher privacy budget.

Untargeted Attack. Albert et al. [2] analyze manipulation attacks in LDP. Compared to Cao et al.’s work, their study mainly focuses on *untargeted* attacks. The attackers aim to skew the original distribution and degrade the overall estimation accuracy of the server.

They suggest for the LDP protocols that the architecture is inherently vulnerable to malicious clients’ manipulations. They suppose a general manipulation attack: the attacker injects M users in N clients in the LDP protocol. These injected users can send arbitrary data sampled from carefully skewed distributions to the server without supposed perturbation. We consider this attacker model corresponds to MGA in [1] and *output-manipulation* (described later) in this paper. We should focus on one of their contributions: they show the general manipulation attack can skew the estimated distribution by $\Omega(\frac{M\sqrt{d}}{\epsilon N})$ in the frequency estimation, which causes more significant error than input-manipulation by about a $\frac{\sqrt{d}}{\epsilon}$ factor (Table 2). The difference is, for example, defined as l_1 -norm of the original and skewed distribution.

Summary. We summarize these notable results in Table 2, showing how effective output-manipulation can attack compared to input-manipulation. The above two

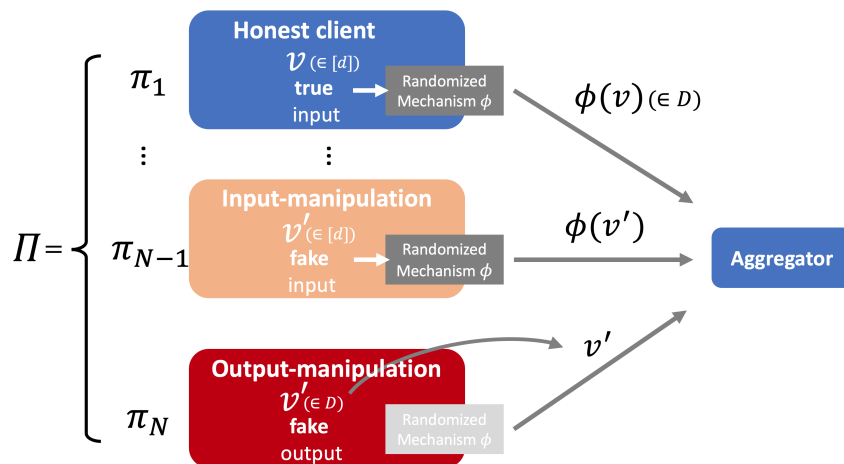


Fig. 1. From top to bottom, normal protocol, input-manipulation attack and output-manipulation attack against an LDP protocol.

previous studies' common conclusion is highlighting the great necessity of enforcing the correctness of users' randomization to defend the output-manipulation attacks.

3 Problem statements

Firstly, we give some notations to LDP protocols, partially following the above-mentioned in Section 2. We denote a single LDP protocol as π_i , where a client c_i sends sensitive data v to server S in ϵ -LDP manner. Encode and perturbation are denoted together as ϕ . ϕ is a probabilistic function (i.e., randomization mechanism) that takes $v \in [d]$ as input and output $y \in D$, such that output space $D = [d]$ if kRR, $D = \{0, 1\}^d$ if OUE, $D = [g]$ if OLH. And we denote overall protocol including all clients as $\Pi = \{\pi_i | i \in [N]\}$.

3.1 Overview of our goal

An attacker against Π injects compromised users into the protocol to send many fake data to a central server. Note that such an attack results in manipulation against a single protocol π by each compromised user. Therefore, we consider security for π , and by protecting security for π , we can naturally protect security for Π . As for the attacker's capability, the attacker can access the implementation of ϕ because this is executed on clients' local, and he knows all parameters and functions including ϕ , ϵ , d , D and $Support(y)$, and employs this information to craft effective malicious outputs. However, in fact, there is little variation in the attacker's behavior because the server can easily deny the protocol if output $y \notin D$. Under such conditions, as shown in Figure 1, we can observe that an attacker can carry out the following two classes of attacks:

Input-manipulation supposes that the attacker can only select input data $v \in [d]$ and cannot interfere with other parameters and functions in π (middle in Figure 1). In other words, the attacker must send $y = \phi(v)$. But in a realistic setting, we should consider it a too strong assumption, as it allows an attacker to have complete control over the local system environment. For example, in targeted attack, RIA corresponds to this class of attacks.

Output-manipulation supposes the attacker can send arbitrary outputs to the server (bottom in Figure 1). This corresponds to the attacker can ignore all parameters and functions ϵ, ϕ by manipulating outputs directly. This attack is an entirely reasonable attack against a general LDP protocol because the server cannot distinguish between true data or fake data. In targeted attack, it corresponds to MGA or the attacks proposed in [2] for untargeted attack. Generally, this class effectively attacks, as shown in Table 2.

An important observation from Section 2.2 is that input-manipulation is much less effective than output-manipulation. Therefore, the natural direction is to defend against output-manipulation and limit the attack to the range of input-manipulation to achieve secure LDP protocols. On the other hand, it is hard to prevent input-manipulation completely. These have been studied in the fields of game theory [27] [26] or truth discovery [28], and we leave such a solution as future work.

Overall, our goal is to mitigate attacks against the LDP protocols by completely defending output-manipulation and limiting to input-manipulation. For this purpose, we consider enforcing the correct mechanism ϕ for protocol π . The key idea is to make the protocol verifiable against malicious clients from a server. In the rest of the paper, we refer to this property as *output-manipulation secure*. (It is also expressed simply as *secure*, and we call it as secure LDP protocol.)

Definition 3 (output-manipulation secure). *An LDP protocol π is output-manipulation secure if any malicious client cannot perform output-manipulation and can only perform input-manipulation against π .*

3.2 Security definitions

In this subsection, we clarify what we should achieve for a secure LDP protocol. Similar to [4], security definitions of secure LDP protocol are consistent with a traditional secure two-party computation (2PC) protocol described in [25]. It considers an ideal world where we can employ Trusted Third Party (TTP) to execute arbitrary confidential computations indeed. And we aim to replace the TTP with a real-world implementation of cryptographic protocol $\pi = (c, S)$ between client c and server S . The protocol's flow when using a TTP is very simple. The client c sends input v to the TTP, and the TTP provides $y = \phi(v)$ to S . After all, c and S never receive any other information; S does not know v , and c does not know y . (S can estimate v from y and ϕ , but c 's privacy should be guaranteed by LDP.)

While it is seemingly apparent that this ideal world's protocols will satisfy our requirements, let us review possible attacks closely. Goldreich [25] summarizes that there are just three types of attacks in malicious model 2PC against

ideal world protocols; (1) denial of participation in the protocol; (2) fake input, not the true one; (3) aborting the protocol prematurely. We cannot hope to avoid these, but (1) and (3) cannot influence the estimation of original data distributions in LDP protocols. (2) is exactly input-manipulation described in the previous subsection. Thus, it is sufficient that the ideal world in 2PC is *output-manipulation secure* (see Def. 3) in the LDP protocols.

Considering the substituted cryptographic protocol $\pi = (c, S)$, let client c as prover \mathcal{P} and server S as verifier \mathcal{V} and $\pi = (\mathcal{P}, \mathcal{V})$. More specifically, we should guarantee secure LDP under the worst case that both \mathcal{P} and \mathcal{V} behave maliciously. The case where \mathcal{P} is malicious is obvious, considering output-manipulation. Still, in the \mathcal{V} 's case, it is because, given the original scenario of LDP, we need to guarantee the privacy of \mathcal{P} . And, we assume \mathcal{P} is a polynomial computational adversary and \mathcal{V} is unbounded.

Following [4], the ideal world protocol can be substituted with protocol π if (a) for any prover algorithm \mathcal{P}^* , \mathcal{V} who receives $\phi(v) = y$ accepts only when \mathcal{P}^* 's secret input is surely v , or otherwise halts with negligible error; (b) for any prover algorithm \mathcal{P}^* , y is indistinguishable from other categories; (c) for any verifier algorithm \mathcal{V}^* , v is indistinguishable from other categories. Additionally, we need to verify that the randomization function ϕ used in the protocol does indeed satisfy ϵ -LDP.

Let $\text{view}_{\mathcal{P}}$ (resp. $\text{view}_{\mathcal{V}}$) as the set of messages generated by the protocol that \mathcal{P} (resp. \mathcal{V}) can observe. And let k as a security parameter that increases logarithmically with cryptographic strength. Then, our security definitions are reduced as following three properties:

- **Verifiability:** This property corresponds to the above-mentioned (a). We consider the protocol is verifiable if it satisfies as follows;

$$\Pr[\mathcal{V} \text{ does not halts } | y \leftarrow \phi(*)] = 1 \text{ and,} \quad (6)$$

$$1 - \Pr[\mathcal{V} \text{ halts } | y \leftarrow \mathcal{P}^*] < \text{negl}(k) \quad (7)$$

where $\text{negl}(k)$ is negligible function in k , $y \leftarrow \phi(*)$ means y is obtained by correct execution of ϕ and $y \leftarrow \mathcal{P}^*$ means y is obtained by \mathcal{P}^* other than correct $\phi(v)$.

- **Indistinguishability:** This property corresponds to (b) and (c). (b) satisfies if $\text{view}_{\mathcal{P}^*}$ has indistinguishable distributions for any input category $v \in [d]$. Formally, we define this property as follows; for any adversary \mathcal{P}^* ,

$$|\Pr[\mathcal{P}^*(\text{view}_{\mathcal{P}^*}, v) = y] - \Pr[\mathcal{P}^*(v) = y]| < \text{negl}(k) \quad (8)$$

where $\text{negl}(k)$ is negligible function in k . This means that a malicious client can use any information obtained from the protocol but only get negligible information about the final output of the server side. Similarly, (c) satisfies if, for any unbounded adversary \mathcal{V}^* ,

$$|\Pr[\mathcal{V}^*(\text{view}_{\mathcal{V}^*}, y) = v] - \Pr[v|y]| < \text{negl}(k) \quad (9)$$

- **Local Differential Privacy:** The randomization mechanism ϕ in the given protocol must satisfy ϵ -LDP as shown in Def. 1. The verification of the correct execution is performed in Eq. (6).

4 Proposed Method

We design secure LDP protocols for kRR, OUE, and OLH, respectively, so that we defend output-manipulations completely. In our method, a major building block is Cryptographic Randomized Response Technique (CRRT) [4] which employs Pedersen’s commitment scheme [21] for secure verifiability using additive homomorphic property, and Naor-Pinkas 1-out-of- n Oblivious Transfer (OT) technique [22] for tricks for a verifiable randomization mechanism. Overall, the proof of validity is based on disjunctive proof [23]. It is a lightweight interactive proof protocol based on a secret sharing scheme and, can perform witness-indistinguishable [31] proofs of knowledge (similar to zero-knowledge proofs). Combined with the security of the encryption scheme proposed in [4], it is possible to securely prove that the output value y is obtained by sampling from a probability distribution that satisfies the ϵ -LDP i.e. $y = \phi(v)$. For simplicity, we explain several phases separately in the following protocol description (Protocol 1, 2), but they can be done simultaneously in the actual implementation.

Before explanation of the protocols in detail, we introduce the following cryptographic setting. Assume that p and q are sufficiently large primes such that q divides $p - 1$, \mathcal{Z}_p has a unique subgroup G of order q . q is the shared security parameter between \mathcal{P} and \mathcal{V} . Security parameter k is $k = \log_2 q_{max}$ such that q_{max} is the maximum value of possible q . We select g and h as a public key. They are two generators of G and, their mutual logarithms $\log_g h$ and $\log_h g$ are hard to compute. We use this public key in the following protocols.

4.1 Secure kRR

Protocol 1 shows the details of the secure version of kRR, an extension of CRRT [4] to satisfy LDP for multidimensional data. As a whole, in the setup phase, both \mathcal{P} and \mathcal{V} prepare the same parameters l, n, z from accuracy parameter *width* and privacy budget ϵ by Algorithm 1. l, n, z identify an categorical probability distributions that satisfies LDP and we use it in 1-out-of- n OT for verifiable random sampling. In mechanism phase, \mathcal{P} creates a vector $\boldsymbol{\mu}$ representing the categorical distribution containing n data where each data μ_i corresponds to one of the categories $[d]$. *width* (i.e., n) is the size of the vector and decides a trade-off between accuracy to approximate LDP and overheads caused by the protocol. For proof P2, we use z^{μ_i} instead of μ_i . All z^{μ_i} is encrypted to y_i by an encryption scheme that combines Pedersen’s commitment and OT. Only the μ_σ , where σ is pre-chosen by \mathcal{V} , can be decrypted correctly. Such a trick allows us to surely perform random sampling from vector $\boldsymbol{\mu}$ representing the categorical distribution. In the proof phase, two proofs are verified in the protocol. The first one is a disjunctive proof for each encrypted data y_i belonging to one of the categories $[d]$ (**P1**). The second one also uses a disjunctive proof that the summation of the vector used as categorical distribution in the OT belongs to one of the possible values (**P2**). There are just d possible values for the summation of $\boldsymbol{\mu}$ (4.(a)).

Algorithm 1 DECIDESHAREDPARAMETERS

Input: $\epsilon, width$
1: $i \leftarrow \lfloor \frac{e^\epsilon}{(d-1)+e^\epsilon} \rfloor$ // as an integer
2: **while** $i > 0$ **do**
3: **if** $(width - i)$ divides $(d - 1)$ **then**
4: $g \leftarrow \gcd(i, width, \frac{width-i}{d-1})$
5: $l, n \leftarrow \frac{i}{g}, \frac{width}{g}$
6: **break**
7: **end if**
8: $i \leftarrow i - 1$
9: **end while**
10: $z \leftarrow \max(\lfloor l, \frac{n-l}{d-1} \rfloor) + 1$
Output: l, n, z

Here, we confirm Protocol 1 is secure. From the protocol, prover and verifier get $\text{view}_{\mathcal{P}} = \{g^a, g^b, g^{ab-\sigma-1}, x_i, x\}$ and $\text{view}_{\mathcal{V}} = \{w_i, y_i, \text{com}_i^{(j)}, c_i^{(j)}, h_i^{(j)}, \text{com}_j, c_i, h_i\}$ for all $i \in [n], j \in [d]$ respectively.

Firstly, we consider indistinguishability. Our encryption scheme (e.g., μ_i is encrypted to y_i) is the same as the one presented in [4], which has been shown that a protocol using the scheme is sufficiently indistinguishable for \mathcal{P}^* and \mathcal{V}^* . That is, it is as hard for \mathcal{P}^* to know about the σ , and also hard for \mathcal{V}^* to guess the distribution of μ and input v . Considering the attacker views, for \mathcal{P}^* , calculating σ from $\text{view}_{\mathcal{P}}$ is as hard as the Decisional Diffie Hellman (DDH) problem. And x and x_i are completely random integers. For \mathcal{V}^* , (w_i, y_i) of $\text{view}_{\mathcal{V}}$ is indistinguishable by the security of the cryptographic scheme, and $(\text{com}_i^{(j)}, c_i^{(j)})$ is also indistinguishable because of the secret sharing scheme [23]. Verifiability is satisfied by proofs, P1 and P2. If both P1 and P2 are verified, \mathcal{V} itself selects one value from the verified vector by OT. Then, for any operation by \mathcal{P}^* , \mathcal{V} can confirm the correctness of the protocol. Hence, verifiability entirely depends on the protocol that proves the P1 and P2. We use disjunctive proofs and Eq.(6) and Eq.(7) are respectively satisfied by the completeness and soundness of the disjunctive proofs shown in [23]. Lastly, Algorithm 1 definitely generates l, n such that $\frac{l}{n} \leq \frac{e^\epsilon}{(d-1)+e^\epsilon}$ and $\frac{n-l}{d-1} \geq \frac{1}{(d-1)+e^\epsilon}$. Hence, because random sampling from μ is equivalent to kRR with $p = \frac{l}{n}, q = \frac{n-l}{d-1}$, at least ϵ -LDP is satisfied.

4.2 Secure OUE

We show the secure version of OUE protocol in Protocol 2. Unlike kRR, OUE sends a d -length bit vector where each i -th bit corresponds to that client likely has the item $i \in [d]$. In OUE, mechanism ϕ performs random bit flips with given constant probability independently for each bit. The Bernoulli distributions, which determine the probabilities of each flip, are approximated by a distribution of n -length bit vectors. As in the case of kRR, verifiable random sampling is achieved by a trick using Pedersen's commitment and OT. However, there are d distribution vectors since it needs for each category.

Protocol 1 secure kRR

Client c as prover \mathcal{P} who holds an secret input $v \in [d]$ and server S as verifier \mathcal{V} . ϵ is privacy budget and $width$ is a parameter representing the degree of approximation.

1. Setup phase.

- (a) \mathcal{P} and \mathcal{V} run DECIDESHAREDPARAMETERS($\epsilon, width$) and prepare l, n, z as shown in Algorithm 1. This is an algorithm for approximating integers l, n, z for given ϵ with as little degradation in accuracy as possible while still satisfying privacy protection.
- (b) \mathcal{V} selects $\sigma \in [n]$. And \mathcal{P} prepares a n -length random number vector $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n)$ where for all $1 \leq i \leq n$, $\mu_i \in [d]$, the vector satisfies $\#\{\mu_i | \mu_i \in \boldsymbol{\mu} \text{ and } \mu_i = v\} = l$ and for all $\{v' | v' \in [d] \setminus \{v\}\}$, $\#\{\mu_i | \mu_i = v'\} = \frac{n-l}{d-1}$ where $\#\{\cdot\}$ returns count on a set.

2. Mechanism phase.

- (a) \mathcal{V} picks random $a, b \leftarrow \mathbb{Z}_q$ and sends g^a, g^b and $g^{ab-\sigma+1}$ to \mathcal{P} .
- (b) For all $i \in \{1, \dots, n\}$, \mathcal{P} performs the following subroutine; (1) Generate (r_i, s_i) at random; (2) Compute $w_i \leftarrow g^{r_i}(g^a)^{s_i} = g^{r_i+as_i}$ and $h_i \leftarrow (g^b)^{r_i}(g^{ab-\sigma+1}g^{i-1})^{s_i} = g^{(r_i+as_i)b+(i-\sigma)s_i}$; (3) Encrypt μ_i to y_i as $y_i \leftarrow g^{z^{\mu_i}} h_i$. Then, send (w_i, y_i) to \mathcal{V} .
- (c) \mathcal{V} computes w_σ^b where σ is what \mathcal{V} choose at setup phase, and computes $g^{\mu_\sigma} \leftarrow \frac{y_\sigma}{h w_\sigma^b}$. And then, find μ_σ from the result and g . Thus, \mathcal{V} receives μ_σ as a randomized output from \mathcal{P} .

3. Proof phase for P1.

- (a) For all $j \in [d] \setminus \{\mu_i\}$, for all $i \in \{1, \dots, n\}$, \mathcal{P} generates challenge $c_i^{(j)}$ and response $s_i^{(j)}$ from \mathbb{Z}_q and prepares commitments $com_i^{(j)} \leftarrow h^{s_i^{(j)}} / (y_i / g^{z^j})^{c_i^{(j)}}$. For $\{\mu_i\}$ and for all $i \in \{1, \dots, n\}$, \mathcal{P} generates $w_i \leftarrow \mathbb{Z}_q$ and let $com_i^{(\mu_i)} = h^{w_i}$. Then, send $com_i^{(j)}$ to \mathcal{V} , for all i, j .
- (b) \mathcal{V} picks $x_i \leftarrow \mathbb{Z}_q$ for all $i \in \{1, \dots, n\}$ and sends it to \mathcal{P} .
- (c) For all $i \in \{1, \dots, n\}$, \mathcal{P} computes $c_i^{(\mu_i)} = x_i - \sum_{j \in [d] \setminus \mu_i} c_i^{(j)}$ and $s_i^{(\mu_i)} = v_i c_i^{(\mu_i)} + w_i$. Then, send $c_i^{(j)}$ and $s_i^{(j)}$ for all i, j to \mathcal{V} .
- (d) Finally, \mathcal{V} checks if $h^{s_i^{(j)}} = b(y_i / g^{z^j})^{c_i^{(j)}}$ for all $j \in [d]$ and $x_i = \sum_{j \in [d]} c_i^{(j)}$, for all $i \in 1, \dots, n$. Otherwise halts.

4. Proof phase for P2.

- (a) For all $j \in [d] \setminus \{v\}$, \mathcal{P} generates challenge c_j and response s_j from \mathbb{Z}_q and prepares commitments $com_j \leftarrow h^{s_j} / (\prod_{i \in \{1, \dots, n\}} y_i / g^{z_j})^{c_j}$ where $Z_j = \frac{n-l}{d-1} \left(\sum_{k \in [d] \setminus \{j\}} z^k \right) + lz^j$. And \mathcal{P} generates $w \leftarrow \mathbb{Z}_q$ and let $com_v = h^w$. Then, send com_j to \mathcal{V} , for all $j \in [d]$.
 - (b) \mathcal{V} picks $x \leftarrow \mathbb{Z}_q$ and sends it to \mathcal{P} .
 - (c) \mathcal{P} computes $c_v = x - \sum_{j \in [d] \setminus \{v\}} c_j$ and $s_v = \left(\sum_{i \in 1, \dots, n} v_i \right) c_v + w$. Then, send c_j and s_j for all j to \mathcal{V} .
 - (d) Finally, \mathcal{V} checks if $h^{s_j} = b(\prod_{i \in \{1, \dots, n\}} y_i / g^{z_j})^{c_j}$ for all $j \in [d]$ and $x = \sum_{j \in [d]} c_j$. Otherwise halts.
-

Protocol 2 secure OUE

\mathcal{P} , $v \in [d]$, \mathcal{V} , $width$, ϵ as with Protocol 1.

1. Setup phase.

- (a) \mathcal{P} and \mathcal{V} set l, n as $\lceil \frac{1}{1+\epsilon} \cdot width \rceil$ and $width$ itself respectively.
- (b) \mathcal{V} selects d random numbers $\sigma = \{\sigma_1, \dots, \sigma_d\}$ where $1 \leq \sigma_j \leq n$. \mathcal{P} prepares d n -length random bit vectors $\vec{\mu} = (\mu_1, \dots, \mu_n)$ such that $\mu_j = (\mu_1^{(j)}, \dots, \mu_n^{(j)})$ where all $\mu_i^{(j)} \in \{0, 1\}$, and the vector satisfies $\sum_i \mu_i^{(j)} = n - l$ if $j = v$ and $\sum_i \mu_i^{(j)} = l$ if $j \neq v$.

2. Mechanism phase.

- (a) \mathcal{V} picks random $a_j, b_j \leftarrow \mathbb{Z}_q$ and sends g^{a_j}, g^{b_j} and $g^{a_j b_j - \sigma_j + 1}$ to \mathcal{P} for all $j \in [d]$.
- (b) For all $j \in [d]$ and $i \in \{1, \dots, n\}$, \mathcal{P} performs the following sub-routine; (1) Generate $(r_i^{(j)}, s_i^{(j)})$ at random; (2) Compute $w_i^{(j)} \leftarrow g^{r_i^{(j)}} (g^{a_j})^{s_i^{(j)}} = g^{r_i^{(j)} + a_j s_i^{(j)}}$ and $h_i^{(j)} \leftarrow (g^{b_j})^{r_i^{(j)}} (g^{a_j b_j - \sigma_j + 1} g^{i-1})^{s_i^{(j)}} = g^{(r_i^{(j)} + a_j s_i^{(j)}) b_j + (i - \sigma_j) s_i^{(j)}}$; (3) Encrypt $\mu_i^{(j)}$ to $y_i^{(j)}$ as $y_i^{(j)} \leftarrow g^{\mu_i^{(j)}} h_i^{(j)}$. Then, \mathcal{P} sends all $(w_i^{(j)}, y_i^{(j)})$ to \mathcal{V} .
- (c) For all $j \in [d]$, \mathcal{V} computes $g^{\mu_{\sigma_j}^{(j)}} \leftarrow y_{\sigma_j}^{(j)} / h^{(w_{\sigma_j}^{(j)})^{b_j}}$. And then, find μ_{σ_j} . Thus, \mathcal{V} receives $[\mu_{\sigma_1}, \dots, \mu_{\sigma_d}]$ as a randomized output from \mathcal{P} .

3. Proof phase for P1.

- (a) For all $j \in [d]$, for all $i \in \{1, \dots, n\}$, \mathcal{P} generates challenge $c_{1-\mu_i^{(j)}, i}^{(j)}$ and response $s_{1-\mu_i^{(j)}, i}^{(j)}$ from \mathbb{Z}_q and prepares commitments $com_{1-\mu_i^{(j)}, i}^{(j)} \leftarrow h^{s_{1-\mu_i^{(j)}, i}^{(j)}} / (y_i^{(j)} / g^{\mu_i^{(j)}})^{c_{1-\mu_i^{(j)}, i}^{(j)}}$. Generate $w_i^{(j)} \leftarrow \mathbb{Z}_q$ and compute $com_{(\mu_i^{(j)}, i)}^{(j)} \leftarrow h^{w_i^{(j)}}$. Then, send $com_{\{0,1\}, i}^{(j)}$ to \mathcal{V} , for all i, j .
- (b) \mathcal{V} picks $x_i^{(j)} \leftarrow \mathbb{Z}_q$ for all $j \in [d]$ and $i \in 1, \dots, n$ and sends it to \mathcal{P} .
- (c) For all $j \in [d]$ and $i \in \{1, \dots, n\}$, \mathcal{P} computes $c_{\mu_i^{(j)}, i}^{(j)} = x_i^{(j)} - c_{1-\mu_i^{(j)}, i}^{(j)}$ and $s_{\mu_i^{(j)}, i}^{(j)} = v_i^{(j)} c_{\mu_i^{(j)}, i}^{(j)} + w_i^{(j)}$. Then, send $c_{\{0,1\}, i}^{(j)}$ and $s_{\{0,1\}, i}^{(j)}$ for all i, j to \mathcal{V} .
- (d) Finally, \mathcal{V} checks if $h^{s_{\{0,1\}, i}^{(j)}} = b(y_i^{(j)} / g^{\{0,1\}})^{c_{\{0,1\}, i}^{(j)}}$ and $x_i^{(j)} = c_{0,i}^{(j)} + c_{1,i}^{(j)}$, for all $i \in \{1, \dots, n\}$ and for all $j \in [d]$. Otherwise halts.

4. Proof phase for P2. (Simplified because it is similar to P1.)

- (a) \mathcal{P} generates and sends all $com_{\{p,q\}}^{(j)}$ to \mathcal{V} .
- (b) \mathcal{V} picks $x_j \leftarrow \mathbb{Z}_q$ for all $j \in [d]$ and sends it to \mathcal{P} .
- (c) \mathcal{P} sends $c_{\{p,q\}}^{(j)}$ and $s_{\{p,q\}}^{(j)}$ for all j to \mathcal{V} .
- (d) \mathcal{V} checks if $h^{s_p^{(j)}} = b(\prod_{i \in 1, \dots, n} y_i^{(j)} / g^{n/2})^{c_p^{(j)}}$ and $h^{s_q^{(j)}} = b(\prod_{i \in 1, \dots, n} y_i^{(j)} / g^l)^{c_q^{(j)}}$ and $x_j = c_p^{(j)} + c_q^{(j)}$ for all $j \in [d]$. Otherwise halts.

5. Proof phase for P3.

- (a) \mathcal{P} computes $h_{sum} \leftarrow \sum_{i,j} h_i^{(j)}$ and sends h_{sum} to \mathcal{V} .
 - (b) \mathcal{V} checks if $h^{h_{sum}} g^{n/2 + l(d-1)} = \prod_{i,j} y_i^{(j)}$. Otherwise halts.
-

In addition, each vector’s distribution is one of two types: j -th vector such that secret input $v = j$ or otherwise (i.e., p or q in Eq. (4)). Thus, we perform independent OT and decide 0 or 1 for d categories and finally, get randomized output $[\mu_{\sigma_1}, \dots, \mu_{\sigma_d}]$.

Then, similar to secure kRR, we must show that all Bernoulli distributions represented by d vectors are correct. Specifically, the proofs are that all elements of bit vectors $\vec{\mu}$ are surely a bit (0 or 1) (**P1**) and distribution of the vectors are surely equivalent to either of p or q of Eq. (4) (**P2**). The number of p and q are 1 and $d - 1$ respectively (**P3**). If all these three proofs are verified, we can confirm the OUE protocol is simulated correctly. Like kRR’s proofs, **P1** and **P2** are proved by d disjunctive proofs. **P3** is based on hardness of discrete logarithm problem. \mathcal{P} cannot find h_{sum} in polynomial time without all correct $h_i^{(j)}$ that is used when encrypting $y_i^{(j)}$. While \mathcal{P} has to release h_{sum} , this is information theoretically indistinguishable from \mathcal{V} for each $h_i^{(j)}$ unless $n = d = 1$. Security statements for the secure OUE protocol are similar to secure kRR. For LDP, as we can see 1.(a) in Protocol 2, we set $q = l/n$ such that $\frac{l}{n} \geq \frac{1}{1+e^\epsilon}$.

4.3 Secure OLH

To make OLH output-manipulation secure, basically, we can use Protocol 1 except that it requires sharing of a hash function and using reduced output category space. As a first step, \mathcal{V} generates and sends a seed s to \mathcal{P} to initialize hash function $H_s : v \rightarrow v'$ where $v \in [d]$ and $v' \in [g]$. \mathcal{V} and \mathcal{P} use the same H_s as a hash function. We can apply Protocol 1 to achieve secure OLH by using category set $[g]$ instead of $[d]$ and sensitive input value v is handled as $v' = H_s(v)$. The rest of the steps are almost the same as kRR.

Even if \mathcal{P}^* , who does not use the hash function correctly, participates the protocol, \mathcal{V} can easily detect it if it sends the output of a different output space, i.e. $y \notin [g]$. If attacker does not use a different output space, the attack can only be equivalent to input-manipulation because \mathcal{V} verifies the correctness of the categorical distribution used in random sampling after applying hash function.

5 Evaluation

In this section, we evaluate and analyze the performance of our proposed protocols. The code in Python is available on github⁴.

Experimental setup. We use an HP Z2 SFF G4 Workstation, with 4-core 3.80 GHz Intel Xeon E-2174G CPU (8 threads, with 8MB cache), 64GB RAM. The host OS is Ubuntu 18.04 LTS. The client and server exchange byte data serialized by pickle (from Python standard library) over TCP. We use $\epsilon = 1.0$ and in OLH, set $g = d/2$ as the hashed space instead of $g = \lfloor e^\epsilon + 1 \rfloor$ for demonstration.

Parameter generator. First, we analyze the approximated probability distribution generated by our proposed method. In secure kRR protocol, we approximate the probability distribution where we generate data to satisfy LDP by

⁴ <https://github.com/FumiyukiKato/verifiable-ldp>




Fig. 2. In secure kRR, with a sufficiently large *width*, categorical distribution by Algorithm 1 can accurately approximate the LDP distributions (left and middle). In secure OUE, it is almost exact discrete approximation with relatively small *width*. (right)

Algorithm 1. Figure 2 shows how accurate the algorithm generates discrete distribution for $\epsilon = (0, 5]$ and for $width = \{100, 1000\}$. The red curve represents probability p for the normal mechanism, and the blue one represents the approximated one. When the *width* is small, there is a noticeable loss of accuracy due to approximation. However, with a sufficiently large *width*, the approximated p has a sufficiently small loss. As the *width* increases, the performance degrades, indicating that there is a trade-off between the accuracy of the probability approximation and the performance. This is true not only for kRR but also for OUE and OLH. For secure OUE, in the right-side of Figure 2, we compare probability q because p is constant in OUE. It is almost exact discrete approximation with small *width*. This is due to the difference in the structure of the vectors that form the probability distribution, with OUE having a simpler structure.

Performance. We evaluate performances of our proposed method. Figure 3 shows total bandwidths, caused in communications of the total protocol, of each three methods for different category sizes. Generally, when increasing category size, total bandwidth also increases. While it increases linearly in OUE, there are fluctuations in kRR and OLH. This is because the probability value that Algorithm 1 approximates may have a smaller denominator (i.e., n) by reduction, which can make the distribution vector smaller. Overall, larger *width* generates almost linear increases in bandwidth. And for the same *width*, secure OUE causes larger communication overhead than others. However, as mentioned in the previous paragraph, secure OUE can approximate the probability distribution with high accuracy using smaller *width*. Hence, in particular, when the number of categories is large, secure OUE is considered to be more efficient by using smaller *width*. Figure 3 shows that, comparing kRR with $width = 1000$ and OUE with $width = 100$, many categories require several times more bandwidth. On the other hand, when the discretized probability distribution can be approximated with a small denominator by reduction, kRR and OLH show a very small bandwidth. When comparing kRR and OLH, OLH is smaller overall. This is due to the fact that the output space is reduced by hashing.

Figure 4 shows total execution time from the time the client sends the first request until the entire protocol is completed. Most of the characteristics are similar to those of bandwidth. As the size of the proofs that need to be computed increases, the execution time is also expected to increase. The only difference is




Fig. 3. With the same *width*, the communication costs of kRR and OLH are small. However, OUE can approximate LDP accurately with small widths (Fig. 2).




Fig. 4. The characteristics of runtime is similar to bandwidth. OLH takes a little longer because of the hashing.

OLH, which takes extra time to execute the hash function. However, as the number of categories becomes larger, the influence becomes smaller.

Therefore, the overhead can be minimized by providing a privacy budget for optimal efficiency for kRR and OLH, and by using different methods for different *width*. The overhead is expected to increase as the number of categories increases, but since the limit on the number of categories is determined to some extent by the use of LDP, we do not think this is a major problem.

At the end, impressively, our method is algorithm-only, making it more feasible than alternatives that assume secure hardware [30] or TEE [29]. Nevertheless, overall, we believe the overhead is acceptable. We believe this is due to the fact that we use relatively lightweight OT techniques as a building block.

6 Conclusion

In this paper, we showed how we prevent malicious clients from attacking to LDP protocol. An important observation was the effectiveness of output-manipulation and the importance of protection against it. Our approach was verifiable randomization mechanism satisfying LDP. Data collector can verify the completeness of executing an agreed randomization mechanism for every possibly malicious data provider. Our proposed method was based on only lightweight cryptography. Hence, we believe it has high feasibility and can be implemented in various and practical data collection scenarios.

References

1. Cao, Xiaoyu, Jinyuan Jia, and Neil Zhenqiang Gong. Data poisoning attacks to local differential privacy protocols. arXiv preprint arXiv:1911.02046 (2019).
2. A. Cheu, A. Smith and J. Ullman, Manipulation Attacks in Local Differential Privacy, in 2021 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, US, 2021 pp. 1-18.
3. Narayan, Arjun, et al. Verifiable differential privacy. Proceedings of the Tenth European Conference on Computer Systems. 2015.
4. Ambainis, Andris, Markus Jakobsson, and Helger Lipmaa. Cryptographic randomized response techniques. International Workshop on Public Key Cryptography. Springer, Berlin, Heidelberg, 2004.
5. Evfimievski, Alexandre, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems. 2003.
6. Dwork, Cynthia, et al. Calibrating noise to sensitivity in private data analysis. Theory of cryptography conference. Springer, Berlin, Heidelberg, 2006.
7. Erlingsson, Úlfar, Vasyl Pihur, and Aleksandra Korolova. "Rappor: Randomized aggregatable privacy-preserving ordinal response." Proceedings of the 2014 ACM SIGSAC conference on computer and communications security. 2014.
8. Apple Differential Privacy Team. Learning with privacy at scale. Machine Learning Journal, 2017.
9. Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In NeurIPS, 2017.
10. Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. In NeurIPS, 2014.
11. Tianhao Wang, Jeremiah Blocki, Ninghui Li, and Somesh Jha. Locally differentially private protocols for frequency estimation. In USENIX Security, 2017.
12. EU GDPR. <https://www.eugdpr.institute/>. Accessed 21 Mar 2021
13. Brazil's General Data Protection Law https://iapp.org/media/pdf/resource_center/Brazilian_General_Data_Protection_Law.pdf. Accessed 21 Mar 2021
14. Facebook Cambridge Analytica data scandal (wikipedia) https://en.wikipedia.org/wiki/Facebook-Cambridge_Analytica_data_scandal.
15. Top10 data breaches of 2020. <https://www.securitymagazine.com/articles/94076-the-top-10-data-breaches-of-2020>. Accessed 21 Mar 2021
16. Peter Kairouz, Keith Bonawitz, and Daniel Ramage. Discrete distribution estimation under local privacy. In ICML, 2016.
17. Tianhao Wang, Bolin Ding, Jingren Zhou, Cheng Hong, Zhicong Huang, Ninghui Li, and Somesh Jha. Answering multi-dimensional analytical queries under local differential privacy. In SIGMOD, 2019.
18. Tianhao Wang, Ninghui Li, and Somesh Jha. Locally differentially private frequent itemset mining. In S&P, 2018.
19. Tianhao Wang, Milan Lopuhaä-Zwakenberg, Zitao Li, Boris Skoric, and Ninghui Li. Locally differentially private frequency estimation with consistency. In NDSS, 2020.
20. Gennaro, Rosario, et al. Quadratic span programs and succinct NIZKs without PCPs. Annual International Conference on the Theory and Applications of Cryptographic Techniques. Springer, Berlin, Heidelberg, 2013.

21. Pedersen, Torben Pryds. Non-interactive and information-theoretic secure verifiable secret sharing. Annual international cryptology conference. Springer, Berlin, Heidelberg, 1991.
22. Naor, Moni, and Benny Pinkas. Efficient oblivious transfer protocols. SODA. Vol. 1. 2001.
23. Cramer, Ronald, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. Annual International Cryptology Conference. Springer, Berlin, Heidelberg, 1994.
24. Bellare, Mihir, and phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. Proceedings of the 1st ACM Conference on Computer and Communications Security. 1993.
25. Oded Goldreich. Secure Multi-Party Computation. Final (Incomplete) Draft, October 27 2002.
26. Do, Cuong T., et al. Game theory for cyber security and privacy. ACM Computing Surveys (CSUR) 50.2 (2017): 1-37.
27. Prelec, Dražen. A Bayesian truth serum for subjective data. science 306.5695 (2004): 462-466.
28. Waguih, Dalia Attia, and Laure Berti-Equille. Truth discovery algorithms: An experimental evaluation. arXiv preprint arXiv:1409.6428 (2014).
29. Sabt, Mohamed, Mohammed Achemlal, and Abdelmadjid Bouabdallah. Trusted execution environment: what it is, and what it is not. 2015 IEEE Trustcom/BigDataSE/ISPA. Vol. 1. IEEE, 2015.
30. Costan, Victor, and Srinivas Devadas. Intel SGX Explained. IACR Cryptol. ePrint Arch. 2016.86 (2016): 1-118.
31. Feige, Uriel, and Adi Shamir. Witness indistinguishable and witness hiding protocols. Proceedings of the twenty-second annual ACM symposium on Theory of computing. 1990.
32. Warner, Stanley L. Randomized response: A survey technique for eliminating evasive answer bias. Journal of the American Statistical Association 60.309 (1965): 63-69.