LDP^3: An Extensible and Multi-Threaded Toolkit for Local Differential Privacy Protocols and Post-Processing Methods

Berkay Kemal Balioglu
Department of Computer Engineering
Koç University
Istanbul, Turkiye
bbalioglu23@ku.edu.tr

Alireza Khodaie
Department of Computer Engineering
Koç University
Istanbul, Turkiye
akhodaie22@ku.edu.tr

M. Emre Gursoy
Department of Computer Engineering
Koç University
Istanbul, Turkiye
emregursoy@ku.edu.tr

***Abstract*—Local differential privacy (LDP) has become a prominent notion for privacy-preserving data collection. While numerous LDP protocols and post-processing (PP) methods have been developed, selecting an optimal combination under different privacy budgets and datasets remains a challenge. Moreover, the lack of a comprehensive and extensible LDP benchmarking toolkit raises difficulties in evaluating new protocols and PP methods. To address these concerns, this paper presents LDP^3 (pronounced LDP-Cube), an open-source, extensible, and multi-threaded toolkit for LDP researchers and practitioners. LDP^3 contains implementations of several LDP protocols, PP methods, and utility metrics in a modular and extensible design. Its modular design enables developers to conveniently integrate new protocols and PP methods. Furthermore, its multi-threaded nature enables significant reductions in execution times via parallelization. Experimental evaluations demonstrate that: (i) using LDP^3 to select a good protocol and post-processing method substantially improves utility compared to a bad or random choice, and (ii) the multi-threaded design of LDP^3 brings substantial benefits in terms of efficiency.**

***Index Terms*—local differential privacy, post-processing, data privacy, frequency estimation, privacy toolkit.**

I. Introduction

In recent years, local differential privacy (LDP) has emerged as a popular notion for privacy-preserving data collection. In LDP, each user perturbs their sensitive data on their device before sharing it with a central server. Since the privacy protection step occurs on the user's device through a randomized algorithm, LDP allows for privacy-preserving data analysis when users do not trust the central server (i.e., data collector). Considering that LDP offers privacy to users while enabling the data collector to estimate aggregate statistics, it has become popular in both research and industry [1], [2]. For example, Google developed RAPPOR to analyze users' default browser homepages and search engines, Apple used LDP to identify popular emojis and trending words for typing recommendations, and Microsoft implemented LDP in Windows 10 to collect app usage telemetry [3]–[5].

The popularity of LDP has led to the development of various LDP protocols, such as GRR, BLH, OLH, RAPPOR, OUE, and SS [6]–[8]. In addition, to improve the utility of server-side estimations, several post-processing (PP) methods have been proposed, such as Base-Pos, Norm, Norm-Cut, Norm-Sub, Norm-Mul, Power, and PowerNS [9], [10]. Yet, no standardized criteria exist for selecting an optimal combination of LDP protocol and PP method across all possible ε and datasets. Furthermore, to advance development in LDP, it would be beneficial to have a toolkit where researchers and practitioners can conveniently integrate their new protocols and PP methods, and benchmark them against existing ones. The existence of such a toolkit would also provide guidance for practitioners in terms of which protocols and methods would yield the highest utility in their specific deployments.

To address these needs, in this paper, we propose and develop LDP^3 (pronounced LDP-Cube, stands for: **L**ocal **D**ifferential **P**rivacy with **P**ost **P**rocessing). LDP^3 is an open-source toolkit[1] containing implementations of 6 LDP protocols, 7 post-processing methods, and several utility metrics in a multi-threaded setup. By housing many of the popular protocols and PP methods, LDP^3 aims to offer a holistic resource for researchers and practitioners to integrate and test their newly proposed methods, or to benchmark existing methods on new applications and datasets. The modular design of LDP^3 enables it to be an extensible privacy toolkit, allowing the addition of new protocols, methods, and metrics as needed. Furthermore, the multi-threaded nature of LDP^3 is particularly beneficial, since the execution and benchmarking of multiple protocols and PP methods with many repetitions (to combat LDP's randomness and achieve statistical significance) causes high execution times. The parallelization in LDP^3 helps to reduce execution times substantially.

There are several potential uses of LDP^3 which makes it beneficial to advance LDP research (from a researcher's perspective) and deployment (from a practitioner's perspective). For example:

- Say that a privacy researcher proposes a new LDP protocol or PP method. The researcher can implement their protocol or method into LDP^3 and experimentally com-

[1] https://github.com/alrzakh/LDPcube

pare it with existing methods. Furthermore, for a newly developed protocol, the best-performing PP method can be found by experimenting with LDP^3.
- Say that a practitioner wants to apply LDP to a real-world data collection task. The practitioner has a surrogate data sample, the ε budget, and the utility metric in hand. Using LDP^3, the practitioner can find which LDP protocol and PP method yields the lowest utility loss and apply this combination to the real-world data collection task.
- Existing PP methods have so far only been validated using a limited number of LDP protocols and settings (e.g., only OLH in [10]). A broader and more holistic benchmarking of different protocol and PP method combinations can be performed using LDP^3 to validate (or challenge) previous findings.

The remainder of this paper is organized as follows. In Section II, we provide the necessary LDP background and explain the differences between LDP^3 and other LDP-related libraries and toolkits. In Section III, we describe the design and current implementation of LDP^3, as well as how to use LDP^3 in practice. In Section IV, we experimentally show the benefits and contributions of LDP^3 from two perspectives: (i) using LDP^3 to select a good protocol and PP method combination indeed helps improve utility substantially, and (ii) multi-threaded design of LDP^3 brings substantial benefits in reducing execution times. Section V concludes the paper.

II. Background and Related Work

A. LDP Background and Notation

In a typical LDP setting, there exist multiple clients (users) and a data collector (server). Each user perturbs their data locally on their device before sending it to the server. After perturbed values are collected, the server performs aggregation and estimation to recover population-level statistics. Since each user's data is perturbed using a randomized LDP algorithm, the server cannot infer exact information about any specific user.

Let $\mathcal{U}$ denote the user population and $\mathcal{D}$ denote the domain of users' values. We denote values in the domain by $v \in \mathcal{D}$. For user $u \in \mathcal{U}$, we denote this user's true value by v_u. We denote the true frequency of v by $f(v)$, its estimated frequency after LDP by $\hat{f}(v)$, and its post-processed frequency by $\tilde{f}(v)$. A randomized mechanism ψ is said to satisfy ε-LDP if the following holds.

Definition 1 (ε-LDP)**.** A randomized mechanism ψ satisfies ε-LDP, if and only if for any two values v_1, v_2 in $\mathcal{D}$:

$$\forall y \in Range(\psi): \quad \frac{\Pr[\psi(v_1) = y]}{\Pr[\psi(v_2) = y]} \leq e^{\varepsilon} \tag{1}$$

where $Range(\psi)$ denotes the set of all possible outputs of ψ.

Here, ε is the parameter that determines the strength of privacy protection. It is often called the *privacy parameter* or *privacy budget*. Smaller ε yields stronger privacy.

B. Related Work

The popularity of LDP has led to its application in a variety of contexts and data analysis tasks, such as high-dimensional data collection [11]–[13], heavy hitter identification [14], [15], set-valued data analysis [16], [17], geospatial data analysis [18], [19], and deep learning [20], [21]. LDP protocols and PP methods serve as the fundamental building blocks of many such applications. There have also been a few works which build practical systems or toolkits for the analysis of LDP protocols. However, LDP^3 has some key differences compared to them, which are described below.

Earlier works such as [6] and [10] focus solely on protocols without post-processing, or evaluate post-processing only on a single protocol. The PURE-LDP package was proposed in [8]; however, it contains a subset of the post-processing methods in LDP^3 and it does not support multi-threading. The MULTI-FREQ-LDP package was proposed in [22]; however, the main focus of MULTI-FREQ-LDP is multi-dimensional and longitudinal frequency estimation [12], [23]. Hence, it does not contain post-processing methods. LDPLENS was proposed in [7]; however, it is focused on protocols' adversarial analysis and does not contain post-processing. Overall, LDP^3 is novel in combining many LDP protocols and post-processing methods under a single umbrella in a multi-threaded setup.

III. LDP^3 Design and Implementation

LDP^3 is an open-source toolkit designed to address the need for a systematic and comprehensive tool to evaluate combinations of different LDP protocols and PP methods. Built in a modular and extensible way, LDP^3 allows users to experimentally find the optimal LDP protocol and PP method for a given dataset, utility metric, and privacy budget ε. It consists of the following primary modules (i.e., components):

- The Protocol Module contains implementations of multiple state-of-the-art LDP protocols such as GRR, BLH, OLH, RAPPOR, OUE, and SS. It offers a standardized interface, allowing developers to add new protocols to LDP^3 by implementing the protocols' user-side perturbation and server-side estimation functions.
- The Post-Processing Module contains implementations of PP methods which take as input the estimated frequencies $\hat{f}(v)$ and produce post-processed $\tilde{f}(v)$. In the current implementation of LDP^3, several PP methods are already included. The module is designed in an extensible way so that new PP methods can be added in the future.
- The Utility Measurement Module contains implementations of utility metrics to measure the differences between $f(v)$ and $\hat{f}(v)$, or $f(v)$ and $\tilde{f}(v)$. Examples of currently implemented metrics include ℓ_1 distance, ℓ_2 distance, and KL-divergence. New metrics can be added as desired.
- The Multi-Threading Module aims to address high execution times when running many combinations of protocols and PP methods. It leverages Python's concurrency features to parallelize the data perturbation and aggregation processes, and significantly improves the computational efficiency of large-scale experiments.

- Finally, the Execution Module provides a command-line interface to run experiments, saving detailed results to an output file or displaying aggregate results on the terminal. This module also facilitates the handling of datasets, e.g., reading and writing to files with standard formats (such as txt or csv).

In the rest of this section, we describe the design and current implementation of each module in more detail.

A. Protocol Module

Several LDP protocols have been developed in the literature [6]–[8]. These protocols are often used as building blocks in more complex data analysis tasks and downstream applications. An LDP protocol can be characterized by two main components: (i) user-side encoding and perturbation to satisfy LDP, and (ii) server-side aggregation and estimation to recover population-level statistics. LDP^3 currently contains the implementations of six protocols: GRR, RAPPOR, OUE, BLH, OLH, and SS. New protocols can be added to LDP^3 by implementing two functions: one function for the protocol's user-side encoding and perturbation, and one function for the protocol's server-side aggregation and estimation. The details of the protocols are provided below.

Generalized Randomized Response (GRR). Randomized response is a method originally introduced for survey data collection, and GRR is an extension of this method designed for LDP. It allows for multi-valued domains (i.e., $|\mathcal{D}| \geq 3$) and works for any privacy parameter ε. Given a user's true value v_u, GRR generates a perturbed value $y_u \in \mathcal{D}$ based on the following probabilities:

$$\Pr[\psi(v_u) = y_u] = \begin{cases} p = \frac{e^{\varepsilon}}{e^{\varepsilon}+|\mathcal{D}|-1}, & \text{if } y_u = v_u \\ q = \frac{1}{e^{\varepsilon}+|\mathcal{D}|-1}, & \text{if } y_u \neq v_u \end{cases} \tag{2}$$

Once the server receives perturbed values from all users, it estimates the frequency of a specific value $v \in \mathcal{D}$ by first calculating $\widehat{C}(v)$, which is the number of users who reported $y_u = v$ as their perturbed value. Then, the estimate $\hat{f}(v)$ is found by:

$$\hat{f}(v) = \frac{\widehat{C}(v) - |\mathcal{U}| \cdot q}{(p-q) \cdot |\mathcal{U}|} \tag{3}$$

RAPPOR. The Randomized Aggregatable Privacy-Preserving Ordinal Response (RAPPOR) protocol, introduced by Google, ensures LDP by encoding a user's value into a bitvector and applying randomized perturbation. Although the original RAPPOR protocol uses Bloom filters for encoding, we describe a simpler version of RAPPOR with unary encoding, which is commonly used in the literature. Each user u initializes a bitvector B_u of length $|\mathcal{D}|$, setting all bits to 0 except for the one corresponding to the user's true value v_u: $B_u[v_u] = 1$. The RAPPOR perturbation mechanism then iterates through each bit $i \in [1, |B_u|]$ and keeps or flips the bit with probabilities defined in the following equation:

$$\forall_{i \in [1,|B_u|]}: \ \Pr[B'_u[i] = 1] = \begin{cases} \frac{e^{\varepsilon/2}}{e^{\varepsilon/2}+1}, & \text{if } B_u[i] = 1 \\ \frac{1}{e^{\varepsilon/2}+1}, & \text{if } B_u[i] = 0 \end{cases} \tag{4}$$

The perturbed bitvector B'_u is sent to the server. After receiving perturbed bitvectors from all users, the server calculates $\widehat{C}[v]$, the count of 1's at index v across all received bitvectors:

$$\widehat{C}[v] = \sum_{u \in \mathcal{U}} B'_u[v] \tag{5}$$

Finally, the server computes the estimate $\hat{f}(v)$ using the formula:

$$\hat{f}(v) = \frac{\widehat{C}[v] + |\mathcal{U}| \cdot (\alpha - 1)}{(2\alpha - 1) \cdot |\mathcal{U}|} \tag{6}$$

where α is the bit-keeping probability: $\alpha = \frac{e^{\varepsilon/2}}{e^{\varepsilon/2}+1}$.

Optimized Unary Encoding (OUE). In OUE, after initializing the bitvector B_u in the same way as in RAPPOR, the perturbed bitvector B'_u is determined according to the following probabilities. These probabilities were mathematically derived to minimize the variance of server-side estimation [6], improving protocol utility.

$$\forall_{i \in [1,|B_u|]}: \ \Pr[B'_u[i] = 1] = \begin{cases} \frac{1}{2}, & \text{if } B_u[i] = 1 \\ \frac{1}{e^{\varepsilon}+1}, & \text{if } B_u[i] = 0 \end{cases} \tag{7}$$

The perturbed bitvector B'_u is sent to the server. After receiving perturbed bitvectors from all users, the server calculates the count $\widehat{C}[v]$ in the same way as in Equation 5. The server then computes the estimate $\hat{f}(v)$ using the following formula:

$$\hat{f}(v) = \frac{2 \cdot \left((e^{\varepsilon}+1) \cdot \widehat{C}[v] - |\mathcal{U}|\right)}{(e^{\varepsilon}-1) \cdot |\mathcal{U}|} \tag{8}$$

Binary Local Hashing (BLH). Both RAPPOR and OUE utilize bitvectors of length $|\mathcal{D}|$, which can lead to significant user-side computation costs and user-server communication costs when $|\mathcal{D}|$ is large. BLH addresses these costs by applying hash functions to reduce the domain size. To address potential issues with hash collisions, BLH uses a set of hash functions $\mathcal{H}$, from which each user selects a different function.

Let $\mathcal{H}$ represent a set of hash functions such that each $H \in \mathcal{H}$ maps a value from $\mathcal{D}$ to an integer in the set $\{0, 1\}$, i.e., $H : \mathcal{D} \to \{0, 1\}$. Each user u with true value v_u randomly selects a hash function H_u from $\mathcal{H}$ and computes the integer $x_u = H_u(v_u)$. The perturbation step in BLH then takes x_u and produces a perturbed value $x'_u \in \{0, 1\}$ according to the following probabilities:

$$\forall_{i \in \{0,1\}}: \Pr[x'_u = i] = \begin{cases} \frac{e^{\varepsilon}}{e^{\varepsilon}+1} & \text{if } x_u = i \\ \frac{1}{e^{\varepsilon}+1} & \text{if } x_u \neq i \end{cases} \tag{9}$$

The user sends the tuple $\langle H_u, x'_u \rangle$ to the server. Once the server receives tuples from all users $u \in \mathcal{U}$, it estimates the value v by first computing $Sup(v)$, which is the total count of tuples where the condition $x'_u = H_u(v)$ holds. The server then estimates $\hat{f}(v)$ as:

$$\hat{f}(v) = \frac{(e^{\varepsilon}+1) \cdot (2 \cdot Sup(v) - |\mathcal{U}|)}{(e^{\varepsilon}-1) \cdot |\mathcal{U}|} \tag{10}$$

Optimized Local Hashing (OLH). Unlike BLH, OLH uses a non-binary output space for hash functions. It allows the

encoding of a value v into an integer within the range $[0, g-1]$, where $g \geq 2$ is a parameter of the protocol. The motivation behind this modification is to overcome the utility loss that occurs in BLH when binary encoding is not optimal. OLH has been shown to provide significant utility improvements over BLH, particularly when ε and $|\mathcal{U}|$ are large. The default value for g is derived as $g = e^{\varepsilon} + 1$ [6].

Let $\mathcal{H}$ represent a set of hash functions such that each $H \in \mathcal{H}$ maps a value from $\mathcal{D}$ to an integer in the range $[0,\ g-1]$, i.e., $H : \mathcal{D} \to [0,\ g-1]$. Each user randomly selects a hash function H_u from $\mathcal{H}$ and computes the integer $x_u = H_u(v_u)$. The perturbation step in OLH then takes x_u and produces a perturbed value $x'_u \in [0, g-1]$ with the following probabilities:

$$\forall i \in [0, g-1] : \Pr[x'_u = i] = \begin{cases} \frac{e^{\varepsilon}}{e^{\varepsilon}+g-1} & \text{if } x_u = i \\ \frac{1}{e^{\varepsilon}+g-1} & \text{if } x_u \neq i \end{cases} \quad (11)$$

The user sends the tuple $\langle H_u, x'_u \rangle$ to the server. Once the server receives tuples from all users $u \in \mathcal{U}$, it estimates the value v by first computing $Sup(v)$, which is the total count of tuples where the condition $x'_u = H_u(v)$ holds. The server then estimates $\hat{f}(v)$ as:

$$\hat{f}(v) = \frac{(e^{\varepsilon} + g - 1) \cdot (g \cdot Sup(v) - |\mathcal{U}|)}{(e^{\varepsilon} - 1) \cdot (g-1) \cdot |\mathcal{U}|} \quad (12)$$

Subset Selection (SS). In the SS protocol, each user reports a subset Z_u of the domain $\mathcal{D}$ to the server. The size of the subset $k = |Z_u|$ is a crucial parameter of the protocol. The default value for k is defined as $k = \frac{|\mathcal{D}|}{e^{\varepsilon}+1}$. User u initializes subset Z_u as empty. SS adds v_u to Z_u with probability $\frac{k \cdot e^{\varepsilon}}{k \cdot e^{\varepsilon} + |\mathcal{D}| - k}$. The remainder of the subset Z_u is constructed as follows:

- If v_u was added to Z_u in the previous step, then $k-1$ elements are selected uniformly at random without replacement from $\mathcal{D} \setminus \{v_u\}$ and added to Z_u.
- If v_u was not added to Z_u in the previous step, then k elements are selected uniformly at random without replacement from $\mathcal{D} \setminus \{v_u\}$ and added to Z_u.

The user sends the resulting Z_u to the server. The server receives the subsets Z_u from all users $u \in \mathcal{U}$. The server defines the parameters σ_k and θ_k as:

$$\sigma_k = \frac{k \cdot e^{\varepsilon}}{k \cdot e^{\varepsilon} + |\mathcal{D}| - k} \quad (13)$$

$$\theta_k = \frac{(k-1) \cdot k \cdot e^{\varepsilon} + (|\mathcal{D}| - k) \cdot k}{(|\mathcal{D}| - 1) \cdot (k \cdot e^{\varepsilon} + |\mathcal{D}| - k)} \quad (14)$$

To estimate $\hat{f}(v)$, the server computes $Sup(v)$, which is the total number of clients in $\mathcal{U}$ whose reported subset Z_u contains v. Then, the server estimates $\hat{f}(v)$ as:

$$\hat{f}(v) = \frac{Sup(v) - |\mathcal{U}| \cdot \theta_k}{(\sigma_k - \theta_k) \cdot |\mathcal{U}|} \quad (15)$$

B. Post-Processing Module

Post-processing (PP) methods take as input the estimated frequencies under LDP $\hat{f}(v)$ and produce post-processed frequencies $\tilde{f}(v)$, with the aim of achieving consistency and utility improvement [9], [10]. Different PP methods result in varying trade-offs between increase in utility and bias; furthermore, different PP methods may be desirable for different datasets and LDP protocols. To facilitate the utilization and benchmarking of various PP methods, LDP^3 includes implementations of commonly used PP methods from the literature, which are described below. New methods can be added to the toolkit in the future.

Base-Pos: Frequencies must be non-negative by definition; however, due to the randomization in LDP, $\hat{f}(v)$ may be negative for some $v \in \mathcal{D}$. Base-Pos addresses this problem by converting all negative estimations to 0.

$$\forall v \in \mathcal{D} : \tilde{f}(v) = \begin{cases} \hat{f}(v) & \text{if } \hat{f}(v) \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

Norm: The sum of frequencies across all $v \in \mathcal{D}$ should equal 1; however, this may not hold due to the randomization in LDP. Norm addresses this problem by adding a constant σ to each frequency so that the sum will equal 1.

$$\forall v \in \mathcal{D} : \tilde{f}(v) = \hat{f}(v) + \sigma, \quad \text{such that} \quad \sum_{v \in \mathcal{D}} \tilde{f}(v) = 1 \quad (17)$$

Norm-Cut: Norm-Cut converts negative and small positive frequencies to 0, and it also ensures that the sum of frequencies equals 1. That is:

$$\forall v \in \mathcal{D} : \tilde{f}(v) = \begin{cases} 0 & \text{if } \hat{f}(v) \leq \theta \\ \hat{f}(v) & \text{if } \hat{f}(v) > \theta \end{cases} \quad (18)$$

where θ is a threshold value. The value of θ is chosen such that $\sum_{v \in \mathcal{D}} \tilde{f}(v) = 1$ is ensured.

Norm-Sub: Norm-Sub converts negative frequencies to 0. Then, it adds a constant δ to the frequencies to ensure that the sum of frequencies equals 1.

$$\forall v \in \mathcal{D} : \tilde{f}(v) = \begin{cases} 0 & \text{if } \hat{f}(v) < 0 \\ \hat{f}(v) + \delta & \text{if } \hat{f}(v) \geq 0 \end{cases} \quad (19)$$

Here, the value of δ is chosen such that $\sum_{v \in \mathcal{D}} \tilde{f}(v) = 1$ is ensured.

Norm-Mul: Norm-Mul converts negative frequencies to 0. Then, instead of an additive factor, it uses a multiplicative factor to the remaining frequencies so that their sum is 1.

$$\forall v \in \mathcal{D} : \tilde{f}(v) = \begin{cases} 0 & \text{if } \hat{f}(v) < 0 \\ \alpha \hat{f}(v) & \text{if } \hat{f}(v) \geq 0 \end{cases} \quad (20)$$

Here, α is the multiplicative factor. Its value is chosen such that $\sum_{v \in \mathcal{D}} \tilde{f}(v) = 1$ is ensured.

Power: The rationale of Power is that many real-world datasets follow a statistical distribution such as a Gaussian or power law distribution. Therefore, Power fits a distribution to the estimated frequencies and aims to minimize the expected

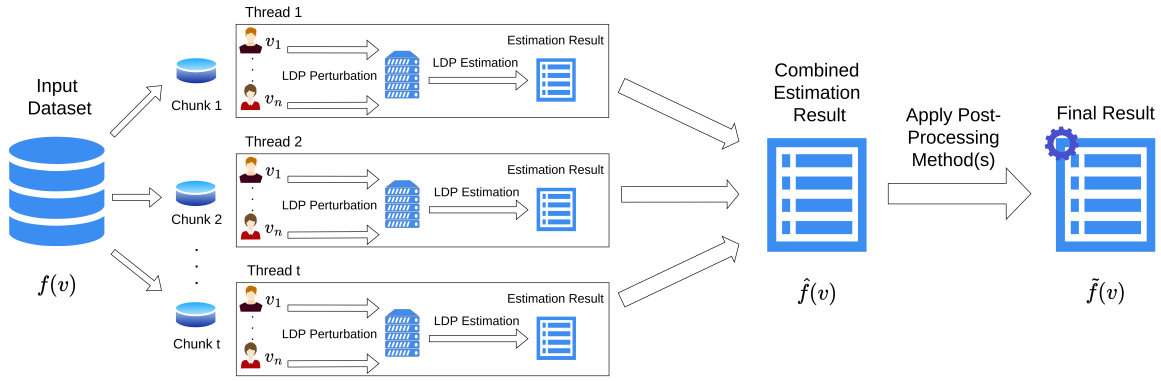

Fig. 1. Multi-threaded execution in LDP^3

square error. For example, with $\tilde{f}(v) \sim P$ where P is a power law distribution, the goal is to minimize:

$$\min_{P} \mathbb{E}\left[\sum_{v\in\mathcal{D}}\left(\hat{f}(v)-\tilde{f}(v)\right)^2\right] \tag{21}$$

PowerNS: PowerNS first applies Power and then uses Norm-Sub on Power's outputs to obtain the final result.

C. Utility Measurement Module

LDP^3 adopts several utility metrics to calculate the errors between frequencies. We exemplify four utility metrics implemented in LDP^3: ℓ_1 distance, ℓ_2 distance, Kullback-Leibler Divergence, and Earth Mover's Distance (EMD). Note that although we present the utility metrics in a way that measures the difference between $f(v)$ and $\tilde{f}(v)$ below, it is also possible to utilize these metrics to measure the difference between $f(v)$ and $\hat{f}(v)$ or the difference between $\hat{f}(v)$ and $\tilde{f}(v)$.

ℓ_1 **distance**, also known as Manhattan distance, measures error using the ℓ_1 norm (absolute value norm). It is defined as:

$$\ell_1 \text{ distance} = \sum_{v\in\mathcal{D}}\left|\tilde{f}(v)-f(v)\right| \tag{22}$$

ℓ_2 **distance**, also known as Euclidean distance, measures error using the ℓ_2 norm (squared error norm). It is defined as:

$$\ell_2 \text{ distance} = \sqrt{\sum_{v\in\mathcal{D}}\left(\tilde{f}(v)-f(v)\right)^2} \tag{23}$$

Kullback-Leibler Divergence (KL-divergence) is a measure of how one probability distribution diverges from a second. Here, $f(v)$ and $\tilde{f}(v)$ are treated as probability distributions:

$$\text{KL-divergence} = \sum_{v\in\mathcal{D}} f(v)\log\left(\frac{f(v)}{\tilde{f}(v)}\right) \tag{24}$$

Earth Mover's Distance (EMD) measures the minimum cost of transforming one distribution into another. It can be interpreted as the amount of work required to transform the original frequency distribution $f(v)$ into the post-processed distribution $\tilde{f}(v)$. Mathematically, it is defined as:

$$\text{EMD} = \min_{\pi} \sum_{v_1,v_2\in\mathcal{D}} \pi(v_1,v_2)\cdot|v_1-v_2| \tag{25}$$

where $\pi(v_1,v_2)$ represents the amount of mass moved from v_1 to v_2, and $|v_1-v_2|$ is the distance between v_1 and v_2.

D. Multi-Threading Module

In practice, finding which protocol and PP method yields the highest utility requires experimenting with many protocol and method combinations. Furthermore, due to the randomized nature of LDP, repeating each experiment multiple times is needed to obtain a reliable and statistically significant result. On the other hand, performing many experiments with many repetitions yields undesirably high execution times.

To address this problem, we implemented LDP^3 in a multi-threaded architecture as shown in Figure 1. The process can be broken down into several key steps. First, the dataset is divided into t equal-sized chunks, where t is the number of threads. Each chunk is assigned to a separate processing thread. A unique random seed is generated for each thread to maintain independence between threads. Second, on each thread, individual user values in the chunk (denoted by v_1 to v_n in Figure 1) undergo LDP perturbation using the selected protocol (e.g., GRR, OLH, RAPPOR, etc.). Intra-thread results are aggregated and estimated within the thread's execution. Each thread executes in parallel and independently from the others. Third, LDP^3 collects estimation results from all threads and combines them via averaging. This results in the estimated frequencies denoted by $\hat{f}(v)$. Finally, $\hat{f}(v)$ undergo selected PP method(s) from the Post-Processing Module (Section III-B). This results in the post-processed frequencies denoted by $\tilde{f}(v)$. The errors between $f(v)$ and $\tilde{f}(v)$ are calculated using metrics implemented in the Utility Measurement Module (Section III-C).

We note two important remarks which are taken into consideration when designing the multi-threading module. First, by default, sizes of all t chunks are equal. This ensures that the averaging performed by LDP^3 is suitable for obtaining $\hat{f}(v)$. If

the sizes of the chunks are different, then weighted averaging (where each chunk is given a weight directly proportional to its size) would ensure correctness. Second, domain size $\mathcal{D}$, privacy budget ε, and other protocol parameters are treated as *global* parameters which are fixed across all threads. This ensures that the protocol behavior remains consistent across all threads.

E. Execution Module

This module provides an interface for executing LDP^3 in practice. The general command structure to run LDP^3 is shown in Figure 2. It can be observed that the command includes several options:

- `-e EPSILON` is used to determine the privacy budget ε.
- `-p PROTOCOLS` specifies the LDP protocol(s) to use. Possible protocols are those implemented in the Protocol Module of LDP^3 (Section III-A). One or more protocols can be used in the same execution. Specifying `"all"` for this option prompts LDP^3 to repeat the experiments with all available protocols in the Protocol Module.
- `-m METHODS` specifies the PP method(s) to use. Possible methods are those implemented in the Post-Processing Module (Section III-B). One or more PP methods can be used in the same execution. Specifying `"all"` for this option prompts LDP^3 to repeat the experiments with all available methods.
- `-r REPEAT`: LDP^3 performs each experiment multiple times (i.e., multiple repetitions) to tackle the inherent randomness of LDP and to achieve statistical significance. The `-r` option is used to determine the number of repetitions per experiment, e.g., 10.
- `-t THREAD_NUMBER` specifies the number of parallel threads to use in the multi-threaded execution of LDP^3.
- `-d DATASET` specifies the dataset path, e.g., `my_dataset.csv`. Currently, it is expected that datasets will contain rows of values where each row corresponds to one user's data. The code in LDP^3 which reads and parses the datasets can be modified to support datasets with different formats.
- `-u UTILITY_METRIC` specifies the utility metric to evaluate errors, selected among those implemented in the Utility Measurement Module (Section III-C).

```
python3 main.py -e EPSILON -p PROTOCOLS
    -m METHODS -r REPEAT -t THREAD_NUMBER
    -d DATASET -u UTILITY_METRIC
```

Fig. 2. Command structure to run LDP^3 with various options and parameters.

IV. Experimental Evaluation

A. Experiment Setup

In this section, we perform experiments to demonstrate that: (i) to improve utility, it is highly beneficial to run experiments with different LDP protocol and PP method combinations using LDP^3 to find a good combination, rather than using a fixed protocol or PP method; and (ii) multi-threading in LDP^3 brings substantial benefits that speed up experiment execution.

We used three real-world datasets for experimentation: BMS-POS, Kosarak, and Porto. We obtained Kosarak from the SPMF Dataset Repository[2], BMS-POS from the public Github repository[3], and Porto is from the ECML-PKDD Taxi Service Prediction Challenge.

- Kosarak contains click-stream data of a Hungarian online news portal. Due to many URLs having very few occurrences (e.g., one or two), we pre-processed the dataset by identifying the top-128 most visited URLs and removed the rest, i.e., $|\mathcal{D}| = 128$. For users who had more than one URL in their resulting stream, the most frequently occurring URL in their stream was picked as their v_u.
- BMS-POS includes market basket sales data from a major electronics retailer, consisting of 515,596 transactions and 1,657 unique items sold. In our experiments, we pre-processed the dataset in a similar fashion to Kosarak by keeping only the top 256 most frequently purchased items.
- Porto contains trips of 442 taxis driving in the city of Porto. While the dataset contains full taxi trips, we pre-processed it by keeping only the starting location of each trip and applied a 15×15 grid for discretization. Each trip was treated as a new v_u. Consequently, we have $|\mathcal{D}|$ = 225 and $|\mathcal{U}|$ = 1,620,157.

Each experiment was conducted 10 times on an Intel Alder Lake Core i7 1255U CPU and the average results are reported. We use ℓ_1 distance as the utility metric.

B. Utility Benefits of LDP^3

In the first set of experiments, we demonstrate the utility benefits of LDP^3 by reporting the errors in frequency estimations without post-processing (w/o PP), with each post-processing method from Section III-B, and the average of all post-processing methods. Results with the Kosarak dataset are shown in Table I, BMS-POS dataset are shown in Table II, and Porto dataset are shown in Table III. For each protocol (each row), the lowest error among all post-processing methods is underlined.

We observe that in many cases, errors are high when there is no post-processing. Post-processing methods consistently help in reducing errors. This consistent improvement demonstrates the effectiveness of post-processing methods and motivates the utility of LDP^3 in enabling the convenient integration and execution of various post-processing methods. Additionally, different post-processing methods yield different amounts of improvement, and the optimal post-processing method varies depending on the protocol and dataset. In other words, there is no post-processing method that performs universally best across all protocols and datasets. For example, in Table II, Norm-Mul is best for GRR, Norm-Sub is best for OLH,

[2]www.philippe-fournier-viger.com/spmf/index.php?link=datasets.php

[3]https://github.com/cpearce/HARM/blob/master/datasets/BMS-POS.csv

TABLE I

ℓ_1 DISTANCES OF ESTIMATIONS UNDER DIFFERENT COMBINATIONS OF LDP PROTOCOLS AND POST-PROCESSING METHODS. KOSARAK DATASET AND ε = 1 ARE USED. ALL VALUES IN THE TABLE ARE $\times 10^{-3}$.

Protocol	w/o PP	Avg. w/ PP	Base-Pos	Norm	Norm-Cut	Norm-Mul	Norm-Sub	Power	PowerNS
GRR	5.65	3.83	3.87	5.65	3.53	3.98	3.08	3.66	3.02
OLH	1.56	1.44	1.33	1.56	1.47	1.39	1.27	1.59	1.49
BLH	1.84	1.65	1.48	1.83	1.71	1.58	1.39	1.84	1.70
OUE	1.48	1.35	1.23	1.47	1.35	1.27	1.19	1.50	1.46
RAPPOR	1.66	1.51	1.37	1.65	1.52	1.47	1.30	1.71	1.61
SS	1.60	1.39	1.34	1.60	1.49	1.43	1.28	1.32	1.21

TABLE II

ℓ_1 DISTANCES OF ESTIMATIONS UNDER DIFFERENT COMBINATIONS OF LDP PROTOCOLS AND POST-PROCESSING METHODS. BMS-POS DATASET AND ε = 1 ARE USED. ALL VALUES IN THE TABLE ARE $\times 10^{-3}$.

Protocol	w/o PP	Avg. w/ PP	Base-Pos	Norm	Norm-Cut	Norm-Mul	Norm-Sub	Power	PowerNS
GRR	10.79	5.94	6.74	10.79	5.15	4.01	4.32	6.36	4.20
OLH	2.14	2.00	1.76	2.14	1.98	1.67	1.66	2.52	2.26
BLH	2.45	2.24	1.95	2.45	2.19	1.83	1.84	2.91	2.53
OUE	2.17	2.03	1.78	2.17	1.98	1.71	1.71	2.53	2.30
RAPPOR	2.21	2.02	1.81	2.21	2.01	1.73	1.73	2.44	2.23
SS	2.08	1.79	1.70	2.08	1.92	1.62	1.61	1.74	1.57

TABLE III

ℓ_1 DISTANCES OF ESTIMATIONS UNDER DIFFERENT COMBINATIONS OF LDP PROTOCOLS AND POST-PROCESSING METHODS. PORTO DATASET AND ε = 1 ARE USED. ALL VALUES IN THE TABLE ARE $\times 10^{-3}$.

Protocol	w/o PP	Avg. w/ PP	Base-Pos	Norm	Norm-Cut	Norm-Mul	Norm-Sub	Power	PowerNS
GRR	5.66	3.97	3.84	5.66	3.45	3.19	3.03	4.33	4.33
OLH	1.17	1.06	0.93	1.17	0.93	0.94	0.87	1.27	1.27
BLH	1.34	1.19	1.05	1.34	1.04	1.06	0.97	1.45	1.45
OUE	1.21	1.09	0.96	1.21	0.97	0.97	0.90	1.33	1.33
RAPPOR	1.27	1.14	1.01	1.27	1.01	1.01	0.94	1.39	1.39
SS	1.18	0.94	0.94	1.18	0.94	0.95	0.88	0.86	0.86

and PowerNS is best for SS. Furthermore, choosing the best protocol and post-processing method is indeed important since the best choice (e.g., SS - PowerNS combination) yields substantially lower error (1.57) compared to a bad choice or an average choice (e.g., $>$ 2.0 error). By offering a range of post-processing methods, LDP^3 can empower researchers and practitioners to experiment with multiple configurations, enabling them to find the best combination of LDP protocol and post-processing method for their specific requirements.

C. Benefits of Multi-Threading

In the second set of experiments, we demonstrate the benefits of the multi-threaded design of LDP^3 by evaluating the impact of multi-threading on the execution times of various protocols. Figures 3, 4, and 5 show the total time required to run each experiment (10 repetitions) on the Kosarak, BMS-POS, and Porto datasets, respectively.

As expected, we observe that the execution times decrease as we increase the number of threads from 1 to 8. However, the decrease is not linear. This is because, as shown in Figure 1, LDP^3 needs to divide the input dataset into t chunks at the beginning. In addition, it needs to combine the estimation results from all threads (chunks) and apply post-processing to the combined estimation result. These actions cannot be multi-threaded; therefore, their time cost cannot be eliminated or reduced by increasing the number of threads. On the other hand, for many protocols, the total execution times are halved or reduced to one-third as we go from a single thread to 4 or more threads. The benefits of multi-threading are especially noticeable in cases where the protocol's execution times are generally high (such as OLH, BLH). Overall, these results highlight the benefits of LDP^3's multi-threading and its suitability for large-scale experimentation.

V. CONCLUSION

In this paper, we introduced LDP^3, an extensible, open-source, and modular toolkit designed to advance research and practical applications of local differential privacy (LDP). By integrating widely used LDP protocols, post-processing methods, and utility metrics within a multi-threaded framework, LDP^3 provides researchers and practitioners with a powerful resource for benchmarking and testing existing or newly proposed methods, as well as exploring the optimized combinations protocols and post-processing methods suitable for their task. Our experimental results highlight the significant utility and efficiency gains enabled by LDP^3. We hope that LDP^3 will be helpful in accelerating LDP research and deploy-

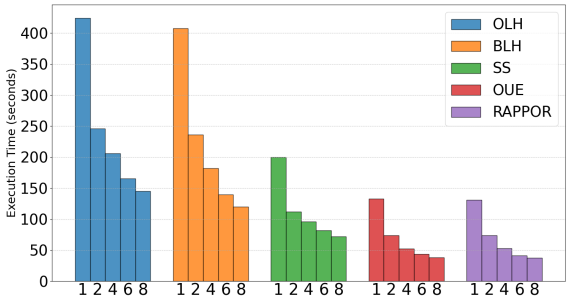

Fig. 3. Execution time vs number of threads (Kosarak).

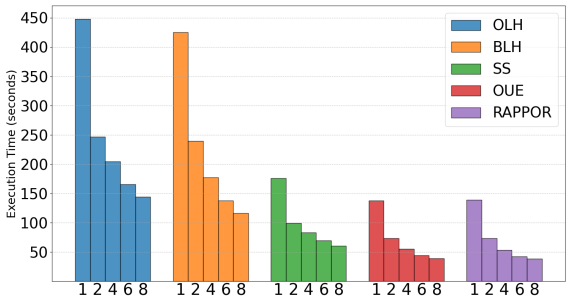

Fig. 4. Execution time vs number of threads (BMS-POS).

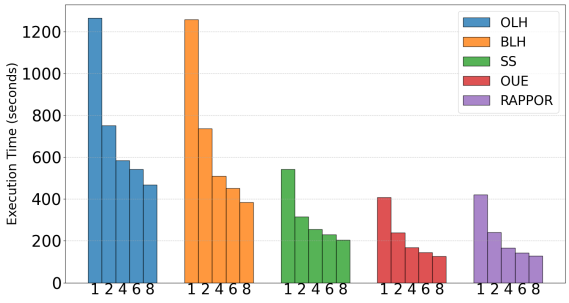

Fig. 5. Execution time vs number of threads (Porto).

ment, as well as fostering collaboration and reproducibility through the public open-source repository.

Acknowledgment

This study was supported by Scientific and Technological Research Council of Türkiye (TUBITAK) under Grant Number 123E179. The authors thank TUBITAK for their support.

References

[1] G. Cormode, S. Jha, T. Kulkarni, N. Li, D. Srivastava, and T. Wang, "Privacy at scale: Local differential privacy in practice," in *Proceedings of the 2018 International Conference on Management of Data*, 2018, pp. 1655–1658.

[2] M. Yang, T. Guo, T. Zhu, I. Tjuawinata, J. Zhao, and K.-Y. Lam, "Local differential privacy and its applications: A comprehensive survey," *Computer Standards & Interfaces*, p. 103827, 2023.

[3] A. G. Thakurta, A. H. Vyrros, U. S. Vaishampayan, G. Kapoor, J. Freudinger, V. V. Prakash, A. Legendre, and S. Duplinsky, "Emoji frequency detection and deep link frequency," Dec. 14 2017, uS Patent App. 15/640,266.

[4] B. Ding, J. Kulkarni, and S. Yekhanin, "Collecting telemetry data privately," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[5] Ú. Erlingsson, V. Pihur, and A. Korolova, "Rappor: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014, pp. 1054–1067.

[6] T. Wang, J. Blocki, N. Li, and S. Jha, "Locally differentially private protocols for frequency estimation," in *26th USENIX Security Symposium*, 2017, pp. 729–745.

[7] M. E. Gursoy, L. Liu, K.-H. Chow, S. Truex, and W. Wei, "An adversarial approach to protocol analysis and selection in local differential privacy," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 1785–1799, 2022.

[8] G. Cormode, S. Maddock, and C. Maple, "Frequency estimation under local differential privacy," *Proceedings of the VLDB Endowment*, vol. 14, no. 11, pp. 2046–2058, 2021.

[9] J. Jia and N. Z. Gong, "Calibrate: Frequency estimation and heavy hitter identification with local differential privacy via incorporating prior knowledge," in *IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2019, pp. 2008–2016.

[10] T. Wang, M. Lopuhaa-Zwakenberg, Z. Li, B. Skoric, and N. Li, "Locally differentially private frequency estimation with consistency," in *Proceedings of the Network and Distributed System Security (NDSS) Symposium*, 2020.

[11] J. S. da Costa Filho and J. C. Machado, "Felip: A local differentially private approach to frequency estimation on multidimensional datasets." in *EDBT*, 2023, pp. 671–683.

[12] H. H. Arcolezi, J.-F. Couchot, B. Al Bouna, and X. Xiao, "Random sampling plus fake data: Multidimensional frequency estimates with local differential privacy," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 47–57.

[13] Z. Zhang, T. Wang, N. Li, S. He, and J. Chen, "Calm: Consistent adaptive local marginal for marginal release under local differential privacy," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 212–229.

[14] Y. Zhu, Y. Cao, Q. Xue, Q. Wu, and Y. Zhang, "Heavy hitter identification over large-domain set-valued data with local differential privacy," *IEEE Transactions on Information Forensics and Security*, 2023.

[15] T. Wang, N. Li, and S. Jha, "Locally differentially private heavy hitter identification," *IEEE Transactions on Dependable and Secure Computing*, vol. 18, no. 02, pp. 982–993, 2021.

[16] S. Wang, Y. Li, Y. Zhong, K. Chen, X. Wang, Z. Zhou, F. Peng, Y. Qian, J. Du, and W. Yang, "Locally private set-valued data analyses: Distribution and heavy hitters estimation," *IEEE Transactions on Mobile Computing*, 2023.

[17] Y. Huang, K. Xue, B. Zhu, D. S. Wei, Q. Sun, and J. Lu, "Joint distribution analysis for set-valued data with local differential privacy," *IEEE Transactions on Information Forensics and Security*, 2024.

[18] D. Hong, W. Jung, and K. Shim, "Collecting geospatial data under local differential privacy with improving frequency estimation," *IEEE Transactions on Knowledge and Data Engineering*, 2022.

[19] Y. Du, Y. Hu, Z. Zhang, Z. Fang, L. Chen, B. Zheng, and Y. Gao, "Ldptrace: Locally differentially private trajectory synthesis," *Proceedings of the VLDB Endowment*, vol. 16, no. 8, pp. 1897–1909, 2023.

[20] B. Wang, Y. Chen, H. Jiang, and Z. Zhao, "Ppefl: Privacy-preserving edge federated learning with local differential privacy," *IEEE Internet of Things Journal*, vol. 10, no. 17, pp. 15 488–15 500, 2023.

[21] S. Truex, L. Liu, K.-H. Chow, M. E. Gursoy, and W. Wei, "Ldp-fed: Federated learning with local differential privacy," in *Proceedings of the Third ACM International Workshop on Edge Systems, Analytics and Networking*, 2020, pp. 61–66.

[22] H. H. Arcolezi, J.-F. Couchot, S. Gambs, C. Palamidessi, and M. Zolfaghari, "Multi-freq-LDPy: Multiple frequency estimation under local differential privacy in python," in *Computer Security – ESORICS 2022*. Springer Nature Switzerland, 2022, pp. 770–775. [Online]. Available: https://doi.org/10.1007/978-3-031-17143-7_40

[23] H. H. Arcolezi, J.-F. Couchot, B. Al Bouna, and X. Xiao, "Improving the utility of locally differentially private protocols for longitudinal and multidimensional frequency estimates," *Digital Communications and Networks*, 2022.