

Locally Differentially Private Sparse Vector Aggregation

Mingxun Zhou*, Tianhao Wang^{†§}, T-H. Hubert Chan[‡], Giulia Fanti*, Elaine Shi*

*Carnegie Mellon University, [†]University of Virginia, [‡]The University of Hong Kong
 {mingxunz,gfanti}@andrew.cmu.edu, tianhao@virginia.edu, hubert@cs.hku.hk, runting@gmail.com

Abstract—Vector mean estimation is a central primitive in federated analytics. In vector mean estimation, each user $i \in [n]$ holds a real-valued vector $v_i \in [-1, 1]^d$, and a server wants to estimate the mean of all n vectors; we would additionally like to protect each user’s privacy. In this paper, we consider the k -sparse version of the vector mean estimation problem. That is, suppose each user’s vector has at most k non-zero coordinates in its d -dimensional vector, and moreover, $k \ll d$. In practice, since the universe size d can be very large (e.g., the space of all possible URLs), we would like the per-user communication to be succinct, i.e., independent of or (poly-)logarithmic in the universe size.

In this paper, we show matching upper- and lower-bounds for the k -sparse vector mean estimation problem under local differential privacy (LDP). Specifically, we construct new mechanisms that achieve asymptotically optimal error as well as succinct communication, either under user-level-LDP or event-level-LDP. We implement our algorithms and evaluate them on synthetic and real-world datasets. Our experiments show that we can often achieve one or two orders of magnitude reduction in error compared with prior work under typical choices of parameters, while incurring insignificant communication cost.

I. INTRODUCTION

Federated analytics and learning allow a cloud provider to learn useful statistics and train machine learning models using data aggregated from a large number of users (e.g., browsing history, shopping records, movie ratings). Since many of these data types are privacy-sensitive, a line of recent work has focused on enabling privacy-preserving federated analytics [8], [10], [27], [28], [32], [54]. A central primitive in privacy-preserving federated analytics is called *vector mean estimation*. Suppose n users each have a real-valued vector $v_i \in [-1, +1]^d$, and the collection of all users’ vectors is called the input configuration, henceforth denoted $\mathbf{v} = (v_1, \dots, v_n) \in [-1, 1]^{d \cdot n}$. The server wants to estimate the mean of the users’ vectors without compromising each individual user’s privacy. *Frequency estimation* [8] can be viewed as a special case of vector mean estimation where each user has a binary vector $v_i \in \{0, 1\}^d$ indicating whether the user owns each of the d items in some universe; the server wants to estimate the frequency of each item. Besides frequency estimation, vector mean estimation is a key building block in numerous applications, such as frequent item mining [50], key-value data aggregation [30], linear regression [43], federated learning model updates [38], and (stochastic) gradient descent [1], [17], [48]. Many of these applications are being

deployed by companies such as Google [27], Apple [51], and Microsoft [22].

In this context, a standard privacy notion is *local differential privacy (LDP)* [34]. Informally, LDP (Def. 7) requires that if a single user changes its input, the distribution of the server’s view in the protocol changes very little. In other words, the transcript observed by the server cannot allow the server to accurately infer any single individual’s input. Two commonly-studied notions of LDP include *user-level LDP* (Def. 10) and *event-level LDP* (Def. 8). In event-level LDP, we want that the distribution of the server’s view to be close under two neighboring input configurations $\mathbf{v} \in [-1, 1]^{d \cdot n}$ and $\mathbf{v}' \in [-1, 1]^{d \cdot n}$ that differ in exactly one coordinate (which may correspond to a single event for a single user). In user-level LDP, we want that the distribution of the server’s view to be close under two neighboring input configurations \mathbf{v} and \mathbf{v}' that differ in the contribution of a single user, possibly involving all d coordinates of that particular user. Unless otherwise noted, throughout this paper, we consider an information-theoretic notion of privacy, i.e., privacy should hold without relying on any computational assumptions.

Sparsity in vector mean estimation. In numerous practical applications, each user’s vector $v_i \in [-1, 1]^d$ is sparse. We say that a vector $v_i \in [-1, 1]^d$ is k -sparse iff at most k coordinates are non-zero. Different users could have different sparsity patterns, i.e., different sets of non-zero coordinates. For example, imagine that the universe consists of the URLs of all websites in the world, and each user’s vector denotes whether a user has visited each URL. Another example is from natural language processing: imagine that the universe is all possible bags-of-words of size three, where a user’s input contains all occurrences that appeared in their emails. In such examples, k is much smaller than the universe size d . Sparsity has also been leveraged as an algorithmic technique in the (non-private) federated learning literature. For example, Konecny et al. [35] showed that sparsifying the gradient vectors in federated learning can significantly reduce communication while maintaining accuracy.

In some cases, the universe may even be too large to efficiently enumerate, e.g., the space of all possible URLs. In such cases, instead of writing down an estimate of the entire mean vector $\bar{v} := \frac{1}{n} \sum_{i=1}^n v_i$, we want the server to instead be able query an estimate of \bar{v}_x for any x of its choice (e.g., the fraction of users that has visited a specific URL of interest).

[§]Work done at CMU.

Due to the prevalence of sparse vectors in real-world applications, we ask the following important question:

Can we achieve locally differentially private k -sparse vector mean estimation with efficient communication and small estimation error?

The most naïve approach is to apply the standard randomized response mechanism [6], [25], [56] to each coordinate — henceforth we refer to this approach as “naïve perturbation”. For the case of event-level LDP, naïve perturbation achieves $\tilde{O}\left(\frac{1}{\epsilon\sqrt{n}}\right) L_\infty$ -error with high probability where ϵ is the privacy budget and $\tilde{O}(\cdot)$ hides (poly-)logarithmic factors. While this simple mechanism actually achieves asymptotically optimal error in light of well-known lower bounds [8], it has a high $O(d)$ communication cost.

One interesting question is whether we can achieve succinct communication that is independent or only logarithmically dependent on the universe size d for sparse vectors. Several prior works [2], [3], [7], [8], [12], [18], [19], [27], [53], [54] have explored this question for 1-sparse vectors, i.e., assuming that each user holds exactly one item out of a large universe of d items. The latest results [7], [8], [12], [19], [54] in this line of work showed how to achieve asymptotically optimal error while paying only logarithmic bandwidth — note that in the 1-sparse case, event-level LDP is the same as user-level LDP up to constant factors.

In comparison, the more general case of k -sparsity is less understood, and currently, we do not have matching upper and lower bounds for schemes with succinct (e.g., logarithmic) communication. Although some prior schemes [30], [43] achieve communication that is succinct in the universe size d under even user-level differential privacy, they suffer from $\Omega(\sqrt{d})$ error, thus making them unsuitable for our motivating scenarios where d can be very large. Other works [46], [55] combine sampling and a 1-sparse mechanism: this approach achieves better asymptotic error than [30], [43] while still maintaining succinct communication, but their asymptotic error is a \sqrt{k} or k factor away from optimal, depending on whether we care about user- or event-level LDP.

A. Our Contributions and Results

We give the first locally private constructions for vector mean estimation that achieve *succinct communication* and *optimal error* (up to polylogarithmic factors). Our contributions include new upper- and lower-bounds, as well as an empirical evaluation of our algorithms.

Upper bounds: communication-efficient LDP mechanisms.

We devise new schemes that satisfy ϵ -LDP (either user-level or event-level) with the following desirable properties:

- *Communication efficiency:* Our mechanisms have communication cost that is independent of the universe size d , and depends only on k , i.e., the maximum number of non-zero coordinates per user.
- *Optimal error:* Our schemes satisfy (nearly) optimal error for any ϵ -LDP mechanism.

Specifically, we prove the following theorems. Although not explicitly stated below, all of our upper bounds below assume the existence of a pseudorandom function (PRF) with parameter λ ; however, the PRF is needed only for measure concentration and not needed for privacy.

Theorem 1 (User-level LDP). *There exists an ϵ -user-level-LDP mechanism for the k -sparse mean vector estimation problem that achieves $O(\log k + \lambda)$ per-client communication, and with probability at least $1 - \beta - \text{negl}(\lambda)$, it achieves $O\left(\frac{1}{\epsilon}\sqrt{\frac{k \log(n/\beta) \log(d/\beta)}{n}}\right) L_\infty$ -error. Moreover, the mechanism is non-interactive, i.e., it involves only a single message from each client to the server.*

Theorem 2 (Event-level LDP). *There exists a non-interactive ϵ -event-level-LDP mechanism for the k -sparse mean vector estimation problem that achieves $O(k \log k + \lambda)$ per-client communication, and with probability at least $1 - \beta - \text{negl}(\lambda)$, it achieves $O\left(\frac{1}{\epsilon}\sqrt{\frac{\log(d/\beta)}{n}}\right) L_\infty$ -error.*

Bassily et al. [8] showed that any event-level LDP mechanism for mean estimation (even when $k = 1$) has to suffer from at least $\Omega\left(\frac{1}{\epsilon}\sqrt{\frac{\log d}{n}}\right)$ error. In light of their lower bound, our event-level LDP mechanism achieves optimal error. In fact, it turns out that our user-level LDP mechanism also achieves (nearly) optimal error but to show this we will need to prove a new lower bound as mentioned shortly below.

Table I compares our results with prior works, and show how we achieve asymptotical improvements. Notice that our event-level scheme consumes more bandwidth than the user-level scheme, partly because the optimal error bound for event-level LDP is more stringent than for user-level LDP. It is an open question whether we can further reduce the bandwidth for event-level LDP while still preserving optimality in error¹

Lower bound for k -sparse LDP vector mean estimation.

We extend the proof technique of Bassily and Smith [8] and prove a new lower bound for any user-level LDP mechanism for vector mean aggregation. Our new lower bound almost tightly matches the upper bound in Theorem 1 (up to a logarithmic gap in n), showing our user-level LDP upper bound achieves nearly optimal error.

Theorem 3 ((Informal) Lower bound for user-level LDP). *Any $(\epsilon, o(\frac{\epsilon}{n \log n}))$ user-level LDP mechanism for vector mean aggregation must suffer from at least $\Omega\left(\frac{1}{\epsilon}\sqrt{\frac{k \log(d/k)}{n}}\right)$ error in expectation.*

Empirical evaluation. We implemented our algorithms with

¹Throughout, we assume that the number of queries made by the server into the estimated mean vector is polynomially or subexponentially bounded in the security parameter (denoted λ) of the PRF, depending on whether the PRF has polynomial or subexponential security. In cases where the server does not query the entire universe d , we take L_∞ error over the queries that are actually made.

TABLE I

Comparison with prior work. The “ k -fold repetition of 1-sparse” and “Sampling + 1-sparse” schemes are strawman constructions explained in Section II.

Name	Event-level LDP		User-level LDP	
	Comm. Cost	L_∞ Error	Comm. Cost	L_∞ Error
k -fold repetition of 1-sparse	$O(k \log d)$	$O(\frac{\sqrt{k}}{\epsilon} \sqrt{\frac{\log d}{n}})$	-	-
Sampling + 1-sparse [46], [55]	same as user-level		$O(\log d)$	$O(\frac{k}{\epsilon} \sqrt{\frac{\log d}{n}})$
Naive Perturbation [6]	$O(d)$	$O(\frac{1}{\epsilon} \sqrt{\frac{\log d}{n}})$	$O(d)$	$O(\frac{k}{\epsilon} \sqrt{\frac{\log d}{n}})$
Harmony [43]	same as user-level		$O(\log d)$	$O(\frac{\sqrt{d}}{\epsilon} \sqrt{\frac{\log d}{n}})$
PCKV [30]	same as user-level		$O(\log d)$	$O(\frac{\sqrt{d}}{\epsilon} \sqrt{\frac{\log d}{n}})$
Ours	$O(k \log k + \lambda)$	$O(\frac{1}{\epsilon} \sqrt{\frac{\log d}{n}})$	$O(\log k + \lambda)$	$O(\frac{\sqrt{k \log n}}{\epsilon} \sqrt{\frac{\log d}{n}})$
Lower Bounds	-	$\Omega(\frac{1}{\epsilon} \sqrt{\frac{\log d}{n}})$	-	$\Omega(\frac{\sqrt{k}}{\epsilon} \sqrt{\frac{\log(d/k)}{n}})$

C++². We evaluated our algorithms using both synthetic and real-world datasets. With the synthetic dataset, we could more easily control the parameters k , d , and n , and we could plot the asymptotic behavior of our algorithms. In comparison with prior communication-efficient works, our algorithms achieve a $5.0\times$ reduction in L_∞ error and a $29.6\times$ reduction in mean square error for both event- and user-level LDP, for a typical choice of parameters, e.g., $n = 10^5$, $d = 10^5$, $k = 64$ and $\epsilon = 1.0$. At the same time, our algorithms consume insignificant communication cost. The report size is smaller or up to a few times larger than a TCP/IP packet headers (20 bytes).

We also tested our algorithms on several real-world datasets. Experiment shows a $1.8\times$ to $7.3\times$ reduction in L_∞ error and a $3.1\times$ to roughly $114.3\times$ reduction in mean square error compared to prior schemes.

Additional contributions. Besides the commonly considered user-level and event-level LDP, as a by-product of our upper bound constructions, we come up with a communication-efficient LDP mechanism under a more generalized L -neighboring notion. Two input configurations $\mathbf{v} = (v_1, \dots, v_n) \in [-1, 1]^{d \cdot n}$ and $\mathbf{v}' = (v'_1, \dots, v'_n) \in [-1, 1]^{d \cdot n}$ are said to be L -neighboring, iff they are otherwise identical except for one user’s coordinates v_i and v'_i , and moreover, $\|v_i - v'_i\|_1 \leq L$. Roughly speaking, a mechanism satisfies (ϵ, δ) -LDP for L -neighboring input configurations if the server cannot (ϵ, δ) -distinguish two L -neighboring input configurations (under the standard distance notion of (ϵ, δ) -differential privacy). Note that the commonly known user- and event-level LDP notions are special cases of the above more generalized notion, for $L = 2k$ and $L = 2$, respectively. Therefore, introducing the generalized L -neighboring notion allows us to study user- and event-level LDP under a more unified lens; and indeed we use it as an intermediate stepping stone to get our main results (Theorems 1, 2, and 3). We believe that this generalized L -neighboring notion can be of independent interest in some practical applications. For example, Abadi et al. [1] considered a gradient clipping technique where each

user would clip its gradient vector to a smaller range (thus pruning excessively large or small values) before sending it to the server. This technique allows them to more tightly bound each user’s contribution.

II. TECHNICAL ROADMAP

In this section, we give an informal technical overview of our results.

A. Background: Optimal LDP mechanisms for 1-sparse vectors

Our algorithm for the generalized k -sparse setting is inspired by the Binary Local Hashing (BLH) algorithm [54], which is an LDP mechanism for the 1-sparse binary vector setting. Let the non-zero coordinate index of client i is x_i . Each client samples a hash function $s_i : [d] \rightarrow \{-1, 1\}$ and computes the bit $z_i = s_i(x_i)$. To ensure ϵ -LDP, client i flips z_i with probability $1/(e^\epsilon + 1)$ and submits the potentially flipped bit (denoted by \tilde{z}_i) to the server. For coordinate x , the server first computes how many clients’ bits match the hash value of x , i.e., $t_i = |\{i \in [n] \mid s_i(x) = \tilde{z}_i\}|$. The server then estimates the mean value of coordinate x as $\frac{e^\epsilon + 1}{e^\epsilon - 1} (\frac{2}{n} t_i - 1)$. The BLH algorithm has an L_∞ error of $O\left(\frac{1}{\epsilon} \sqrt{\frac{\log d}{n}}\right)$, which is proven to be asymptotically optimal for 1-sparse vectors [8]. However, we will see next that adapting this approach to the k -sparse setting is nontrivial.

B. Warmup: An event-Level LDP mechanism for frequency estimation

For simplicity, we first focus on the special case of designing a *frequency estimation* mechanism that satisfies *event-level* LDP. Recall that the frequency estimation problem is a special case of our general formulation of vector mean estimation. In frequency estimation, each client $i \in [n]$ has a binary vector $v_i \in \{0, 1\}^d$, denoting whether the client owns each item from a universe of d items. The server wants to estimate the frequency of each item. Once we understand how to design an event-level LDP mechanism for frequency estimation, we can later extend our techniques to 1) support *user-level* LDP; and

²<https://github.com/wuwuz/sparse-vector-aggregation>

2) support the more general case of vector mean estimation where the client’s vector is from a real domain.

Strawman: k -fold repetition of the 1-sparse scheme. Recall that prior works [8], [19], [54] have proposed 1-sparse frequency estimation mechanisms that achieve optimal error, that is, $\tilde{O}(\frac{1}{\epsilon\sqrt{n}})$ error, incurring only logarithmic communication. In our problem, each client owns k items rather than 1. Therefore, a strawman idea is through a k -fold repetition of the 1-sparse scheme. Specifically, each client can pretend to be k virtual clients, and each virtual client owns only one item. Imagine that we run a 1-sparse scheme over these kn virtual clients. Since each client acts as k virtual clients, its communication cost is $\tilde{O}(k)$ which is independent of the universe size d . The resulting L_∞ error would be $\tilde{O}(\frac{1}{\epsilon\sqrt{kn}})$ over all kn virtual clients. In reality, we want to take the mean over the n real clients. After renormalizing, the actual error is $\tilde{O}(\frac{\sqrt{k}}{\epsilon\sqrt{n}})$.

This strawman scheme gives non-trivial bounds, but does not achieve optimal $\tilde{O}(\frac{1}{\epsilon\sqrt{n}})$ error.

Our idea. We devise a new scheme that combines the elegant ideas behind the 1-sparse mechanism by Wang et al. [54] with a new random binning idea. Our approach is as follows:

- *Each client* $i \in [n]$ does the following:
 - 1) Sample two random hash functions $h_i : [d] \rightarrow [k]$ and $s_i : [d] \rightarrow \{-1, 1\}$.
 - 2) Let $\{x_1, \dots, x_k\} \in [d]^k$ denote the k items belonging to client $i \in [n]$. For each $j \in [k]$, place x_j into the hash bin indexed $h_i(x_j)$. Note that in total, there are k hash bins per client.
 - 3) For each hash bin $j \in [k]$, compute $B_{i,j} := \sum_{x \in \text{bin}_j} s_i(x) + \text{Lap}(\frac{1}{\epsilon})$ where $\text{Lap}(\frac{1}{\epsilon})$ denotes Laplacian noise of average magnitude $\frac{1}{\epsilon}$.
 - 4) Send to the server the tuple $(h_i, s_i, \{B_{i,j}\}_{j \in [k]})$ where h_i and s_i denote the description of the two hash functions.
- *Server* estimates the fraction of clients that own an arbitrary item $x^* \in [d]$ using the following expression: $\frac{1}{n} \sum_{i \in [n]} B_{i, h_i(x^*)} \cdot s_i(x^*)$.

As mentioned later, to get our desired bounds, we need the hash functions h_i and s_i to be pseudorandom — however, we stress that the pseudorandomness assumption is needed only for load-balancing among the hash bins and not for proving privacy. In other words, our scheme satisfies information-theoretic LDP. Specifically, to sample a pseudorandom function (PRF), the client samples a random seed whose length is related to the strength of pseudorandomness and independent of d . To send the description of the hash function to the server, the client sends the pseudorandom seed to the server.

Finally, in practice, the client can clip each $B_{i,j}$ to an integer value between $[-\tilde{O}(k), \tilde{O}(k)]$ before sending it to the server — this does not affect the privacy analysis or our asymptotic error bound. In this case, the per-client communication of our scheme is at most $O(k \log k)$ plus the description of the hash function (e.g., the seed of a PRF).

Informal utility analysis. To gain intuition, we present an informal analysis of our scheme. The formal proofs (for the more generalized vector mean estimation scheme) are deferred to Section IV-A. Note that understanding the utility analysis also helps to understand why the scheme works.

Henceforth, we use bin_j^i to denote the set of items client i places into its j -th bin (and when it is clear from the context which client i we are referring to, we may omit i). Let $C_{i,j} = \sum_{x \in \text{bin}_j^i} s_i(x)$ be the true aggregated “count” of the j -th bin belonging to the i -th client. Suppose that the server wants to know the frequency of item $x^* \in [d]$. To do this, the server computes the summation $\sum_{i \in [n]} B_{i, h_i(x^*)} \cdot s_i(x^*) = \sum_{i \in [n]} C_{i, h_i(x^*)} \cdot s_i(x^*) + \sum_{i \in [n]} \text{Lap}(\frac{1}{\epsilon})$ — note that here we have not normalized the sum with the $\frac{1}{n}$ factor yet, we can defer this step to the end. The second part of the summation $\sum_{i \in [n]} \text{Lap}(\frac{1}{\epsilon})$, is the summation of n independent $\text{Lap}(\frac{1}{\epsilon})$ noises, and thus its magnitude is roughly $\tilde{O}(\frac{\sqrt{n}}{\epsilon})$. The first part of the summation $\sum_{i \in [n]} C_{i, h_i(x^*)} \cdot s_i(x^*)$ can be further decomposed into two sources of contributions:

- 1) Each client i who owns x^* contributes one $+1$ term to the summation because $(s_i(x))^2 = 1$.
- 2) For each client i and each item $x \neq x^*$ owned by the client such that $h_i(x) = h_i(x^*)$, it contributes $s_i(x) \cdot s_i(x^*)$ to the summation, which is a random choice of -1 or $+1$ assuming that $s_i(\cdot)$ is a random oracle.

Thus, 1) corresponds to the true count of the item x^* , whereas 2) can be viewed as the result of a random walk of expected length $O(n)$, i.e., a random noise of magnitude roughly $\tilde{O}(\sqrt{n})$. In particular, the length of this random walk is upper bounded by the total load of the hash bins $\sum_{i \in [n]} |\text{bin}_{h_i(x^*)}^i|$, which is $O(n)$ assuming that each h_i is a random oracle.

Summarizing the above, the estimated count is the true count plus roughly $\tilde{O}(\frac{\sqrt{n}}{\epsilon})$ noise. Finally, when the server normalizes the above sum by $\frac{1}{n}$ to compute the average, the resulting error becomes $\tilde{O}(\frac{1}{\epsilon\sqrt{n}})$. Note that in Theorem 2, the precise expression for the error bound has an extra $\log \frac{d}{\beta}$ term which we ignore here, where β is the failure probability for the error bound. Specifically, the $\log d$ term arises from taking a union bound over the universe of d elements and the $\log \frac{1}{\beta}$ term comes from a precise measure concentration bound on the error — we defer these precise calculations to the subsequent technical sections.

At this point, it is helpful to observe that in this construction, the error comes from two sources — this observation will later help us to generalize the scheme to user-level LDP:

- *Noise component:* the first source of error is the sum of n independent $\text{Lap}(\frac{1}{\epsilon})$ noises, one for each $\text{bin}_{h_i(x^*)}^i$ where $i \in [n]$;
- *Colliding items component:* the second source of error is the random contribution of either $+1$ or -1 from each element $x \neq x^*$ that each client i places into its bin $\text{bin}_{h_i(x^*)}^i$.

Remark 4. In the above, we assumed that the hash functions

h_i 's and s_i 's are random oracles. In practice, we instantiate the hash functions using pseudorandom functions.

Informal privacy analysis. We now give an informal privacy analysis, while deferring the formal proofs to Section IV-A. We want to show that the scheme satisfies ϵ -event-level-LDP. Fix the hash functions $h_1, \dots, h_n, s_1, \dots, s_n$, and consider two input configurations $\mathbf{v}, \mathbf{v}' \in \{0, 1\}^{d \cdot n}$ that differ in only one position. Let $C_{i,j} = \sum_{x \in \text{bin}_{i,j}} s_i(x)$ be the true aggregated “count” of the j -th bin belonging to the i -th client, when the input configuration is \mathbf{v} ; and let $C'_{i,j}$ be the corresponding quantity when the input configuration is \mathbf{v}' . It must be that all $C_{i,j}$ and $C'_{i,j}$ are the same everywhere except for one bin j^* corresponding to one client i^* . Moreover, for the only location where they differ, it must be that $|C_{i^*,j^*} - C'_{i^*,j^*}| \leq 1$. Having observed this, it is not too hard to show that adding Laplacian noise of average magnitude $\frac{1}{\epsilon}$ to each bin suffices for achieving ϵ -event-level-LDP.

C. Extension: A user-level LDP mechanism for frequency estimation

One trivial way to obtain user-level LDP is to directly use the aforementioned warmup scheme and simply apply standard privacy composition theorems [26], [33] to reset the parameters. Specifically, to achieve ϵ -user-level-LDP, we would need to plug in a privacy parameter of $\frac{\epsilon}{k}$ when invoking the warmup scheme. This results in an error bound of $\tilde{O}(\frac{k}{\epsilon\sqrt{n}})$ which is an $\tilde{O}(\sqrt{k})$ factor away from optimal.

Strawman: sampling + 1-sparse mechanism. Another strawman idea for each client to randomly sample 1 item out of its k items, apply the 1-sparse mechanism to the sampled items, and finally, renormalize the estimate accordingly [46], [55]. Unfortunately, it is not hard to show that the resulting error would again be $\tilde{O}(\frac{k}{\epsilon\sqrt{n}})$, an $\tilde{O}(\sqrt{k})$ factor away from optimal. Note also that if a client has strictly fewer than k items, it needs to first pad its input to k with filler items, and then apply the the sampling and 1-sparse mechanism.

Our approach. Our approach is to generalize our warmup mechanism. Suppose we want to achieve (ϵ, δ) -LDP under L -neighboring input configurations. Recall that two input configurations $\mathbf{v} = (v_1, \dots, v_n) \in \{0, 1\}^{d \cdot n}$ and $\mathbf{v}' \in (v'_1, \dots, v'_n) \in \{0, 1\}^{d \cdot n}$ are L -neighboring iff they differ in only one user's contribution v_i and v'_i , and moreover, $|v_i - v'_i|_1 \leq L$. Note that user-level LDP is simply a special case where $L = 2k$. In other words, we want the server's view to be (ϵ, δ) -close for two input configurations $\mathbf{v} = (v_1, \dots, v_n) \in \{0, 1\}^{d \cdot n}$ and $\mathbf{v}' \in (v'_1, \dots, v'_n) \in \{0, 1\}^{d \cdot n}$ under the distance notion of the standard (ϵ, δ) -differential privacy definition [25]. In our reasoning below, we will carry around the parameter L , and in the end, we can plug in $L = 2k$ to get the user-level LDP result. However, as noted earlier in Section I, the more general scheme parametrized by L can be of independent interest.

Our generalized scheme is almost the same as the warmup scheme, except with the following modifications:

- Each client now has b hash bins rather than k bins as in the warmup scheme. For now, we leave the choice of b unspecified, and work out the optimal choice later.
- Each client $i \in [n]$ now computes the noisy sum $B_{i,j}$ as $B_{i,j} := \sum_{x \in \text{bin}_j} s_i(x) + \text{Lap}(\frac{1}{\epsilon'})$ where

$$\epsilon' = \frac{O(\epsilon)}{\min(L, \sqrt{bL \log \frac{b}{\delta}})}.$$

Informal privacy analysis. Consider two L -neighboring input configurations \mathbf{v} and \mathbf{v}' , and fix all hash functions h_1, \dots, h_n and s_1, \dots, s_n . Let $C_{i,j} := \sum_{x \in \text{bin}_j} s_i(x)$ be the true “count” of bin_j^i under \mathbf{v} and let $C'_{i,j}$ be the corresponding quantity under \mathbf{v}' . Now, consider the vectors $\mathbf{C} = \{C_{i,j}\}_{i \in [n], j \in [b]}$ and $\mathbf{C}' = \{C'_{i,j}\}_{i \in [n], j \in [b]}$. We want to show that $|\mathbf{C} - \mathbf{C}'|_1 \leq \min(L, O(\sqrt{bL \cdot \log \frac{b}{\delta}}))$ with probability $1 - \delta$. If so, adding the aforementioned noise is sufficient for achieving (ϵ, δ) -LDP under L -neighboring. Now, $|\mathbf{C} - \mathbf{C}'|_1 \leq L$ is easy to see. Therefore, it suffices to show that $|\mathbf{C} - \mathbf{C}'|_1 \leq O(\sqrt{bL \cdot \log \frac{b}{\delta}})$ with probability $1 - \delta$. Due to standard measure concentration bounds, when we change \mathbf{v} to \mathbf{v}' , for any fixed bin_j^i , it holds that $|C_{i,j} - C'_{i,j}| \leq O(\sqrt{\frac{L \cdot \log \frac{b}{\delta}}{b}})$ with probability $1 - \frac{\delta}{b}$. Taking a union bound over all b bins, we have that $|\mathbf{C} - \mathbf{C}'|_1 \leq O(\sqrt{bL \cdot \log \frac{b}{\delta}})$ with probability $1 - \delta$.

Informal utility analysis and optimal choice of b . As in the earlier event-level LDP scheme, the error in the final summation $\sum_{i \in [n]} B_{i, h_i(x^*)} \cdot s_i(x^*)$ — without normalizing it with the $1/n$ factor yet — comes from two sources:

- *Noise component.* The noise component consists of the summation of n independent $\text{Lap}(\frac{1}{\epsilon'})$ noises where $\epsilon' = \frac{O(\epsilon)}{\min(L, \sqrt{bL \log \frac{b}{\delta}})}$. Thus, the total noise is roughly $O(\frac{1}{\epsilon} \cdot \sqrt{n} \cdot \min(L, \sqrt{bL \log \frac{b}{\delta}}))$.
- *Colliding items component.* The contribution from all colliding elements can be viewed as a random walk of length that is equal to the number of colliding elements in all n bins $\{\text{bin}_{h_i(x^*)}^i\}_{i \in [n]}$. The number of colliding elements is concentrated around its expectation $\frac{nk}{b}$ with high probability, and thus the colliding items component results in roughly $\tilde{O}(\sqrt{\frac{nk}{b}})$ error.

The total error is minimized when the noise component is roughly equal to the contribution from colliding elements, and we derive the optimal choice of b as

$$b = \begin{cases} \tilde{\Theta}\left(\frac{k}{L^{\frac{2}{3}}}\right), & \text{if } L \leq k^{\frac{3}{2}} \\ \tilde{\Theta}\left(\sqrt{\frac{k}{L}}\right), & \text{otherwise.} \end{cases} \quad (1)$$

When the above optimal b is chosen correspondingly, both error components are roughly equal (omitting the logarithmic factors). Specifically, for the case when $L \leq k^{\frac{3}{2}}$, both error components are roughly $\tilde{O}\left(\frac{1}{\epsilon} L \sqrt{n}\right)$; for the other case, both error components are $\tilde{O}\left(\frac{1}{\epsilon} (kL)^{\frac{1}{4}} \sqrt{n}\right)$. Keep in mind

that for our final error bound, we need to apply an extra $1/n$ normalizing factor to the above terms.

Summarizing the above, we obtain a mechanism that satisfies (ϵ, δ) -LDP under L -neighboring, with per-client communication cost $\tilde{O}(b)$ where b is shown in Equation (1), and its choice depends on whether $L \leq k^{\frac{1}{3}}$. Further, the scheme achieves $\tilde{O}\left(\frac{1}{\epsilon} \min(L, (kL)^{\frac{1}{4}}) \sqrt{\frac{1}{n}}\right)$ error after applying the extra $1/n$ normalization factor.

For the special case of user-level-LDP, which can be captured by k -neighboring LDP, using the above calculation, we conclude that the optimal choice of b should be $b = 1$. In this case, the error is roughly $\tilde{O}\left(\frac{\sqrt{k}}{\epsilon\sqrt{n}}\right)$. So far, the scheme described above achieves (ϵ, δ) -LDP with a non-zero δ . However, for the special case of user-level LDP, we can use an additional clipping technique to obtain ϵ -LDP. We defer the detailed exposition of this technique to subsequent sections.

Finally, observe that for $L = 1$, i.e., for event-level-LDP, the optimal choice of $b = O(k)$. This shows that our event-level-LDP scheme in Section II-B is also a special case of the above more generalized scheme.

D. Generalizing to real-valued vectors

In the more general case, each client i holds a real-valued vector $v_i \in [-1, 1]^d$, with at most $k \ll d$ non-zero coordinates. For example, each non-zero coordinate may represent the rating a user has given to a movie that it has watched. Chances are, each user has watched relatively few (k) movies out of the entire universe of d movies.

It is not difficult to generalize the aforementioned schemes (Sections II-B and II-C) to real-valued vectors. The only modification is the following: each client i now computes $B_{i,j}$ as follows for $j \in [b]$ where b denotes the number of bins per client:

$$B_{i,j} := \sum_{x \in [d]} s_i(x) \cdot v_{i,x} + \text{Lap}\left(\frac{1}{\epsilon'}\right)$$

where $v_{i,x}$ denotes the x -th coordinate of the client's vector v_i , and the choice of ϵ' is the same as before. Note that since our event-level-LDP scheme (Section II-B) is a special case of the scheme in Section II-C, the above works for the event-level-LDP scheme too.

The proof of the above generalized scheme is similar in spirit to the binary case but requires more careful calculation. In the subsequent technical sections, we directly prove the real-valued case, since this is the more general form.

E. Lower bound

The framework in Bassily and Smith [8] provides a lower bound for the event-level LDP under 1-sparse setting. Our event-level LDP upper bound tightly matches the lower bounds and therefore closes the case for event-level LDP. We observe that it is straightforward to extend Bassily and Smith [8]'s proof to user-level-LDP, and the resulting lower bound matches the error achieved by our earlier upper bound. We defer the detailed presentation of this lower bound to Section VI.

III. PRELIMINARIES AND DEFINITIONS

A. Background on Differential Privacy

Differential privacy was first proposed by Dwork et al. [25] and has since become a widespread privacy notion.

Definition 5 ((ϵ, δ) -close). *We say the distributions of two random variables, X and X' are (ϵ, δ) -close iff they have the same domain D and for every subset $S \subseteq D$,*

$$\Pr[X \in S] \leq e^\epsilon \Pr[X' \in S] + \delta.$$

Definition 6 ((ϵ, δ) -Differential Privacy). *A function f is (ϵ, δ) -DP w.r.t. some neighboring relation \sim on its input domain iff for every pair $v, v' \in \text{Domain}(f)$, s.t. $v \sim v'$, the distributions of $f(v)$ and $f(v')$ are (ϵ, δ) -close.*

If a function f is $(\epsilon, 0)$ -DP, we also say that f is ϵ -DP for short (w.r.t. the neighboring relation \sim).

B. Sparse Vector Mean Estimation

Consider n clients, indexed by the set $[n] = \{1, 2, \dots, n\}$. Each client has a real-value vector $v_i \in [-1, 1]^d$. Also, each vector v_i is k -sparse, i.e., it has at most k non-zero coordinates. Different clients may have different non-zero coordinates. We use the notation $\mathbf{v} := (v_1, \dots, v_n)$ to denote all clients' inputs, and we also refer to \mathbf{v} as an input configuration. A server wants to estimate the mean vector, $\bar{v} = \frac{1}{n} \sum_{i \in [n]} v_i$ through a non-interactive mechanism.

In a non-interactive mechanism, each client sends a **single** message to the server, and the server then computes an estimate of the mean vector $\bar{v} = \frac{1}{n} \sum_{i \in [n]} v_i$. Both the clients and the server can make use of randomness in their computation.

Henceforth, let \sim denote some symmetric neighboring relation defined over two input configurations $\mathbf{v} \in [-1, 1]^{d \cdot n}$ and $\mathbf{v}' \in [-1, 1]^{d \cdot n}$.

Definition 7 (Local differential privacy (LDP)). *A non-interactive mechanism \mathcal{M} satisfies (ϵ, δ) -LDP w.r.t. the neighboring relation \sim , iff for any two input configurations $\mathbf{v} \in [-1, 1]^{d \cdot n}$ and $\mathbf{v}' \in [-1, 1]^{d \cdot n}$ such that $\mathbf{v} \sim \mathbf{v}'$, it holds that*

$$\Pr[\mathbf{view}_{\mathcal{M}}(\mathbf{v}) \in S] \leq e^\epsilon \Pr[\mathbf{view}_{\mathcal{M}}(\mathbf{v}') \in S] + \delta$$

where $\mathbf{view}_{\mathcal{M}}(\mathbf{v})$ is a random variable representing the server's view upon input configuration \mathbf{v} ; in particular, the view consists of all messages received by the server.

If a mechanism satisfies $(\epsilon, 0)$ -LDP, we also say that it satisfies ϵ -LDP (w.r.t. to some neighboring relation \sim).

Definition 8 (Event-level LDP). *We say that a mechanism satisfies (ϵ, δ) -event-level-LDP, iff it satisfies (ϵ, δ) -LDP w.r.t. the following neighboring relationship: two input configurations $\mathbf{v} = (v_1, \dots, v_n) \in [-1, 1]^{d \cdot n}$ and $\mathbf{v}' = (v'_1, \dots, v'_n) \in [-1, 1]^{d \cdot n}$ are considered neighboring, iff they differ in at most one position (i.e., one coordinate contributed by one user).*

Definition 9 (User-level LDP). *We say that a mechanism satisfies (ϵ, δ) -user-level-LDP, iff it satisfies (ϵ, δ) -LDP w.r.t. the following neighboring relationship: two input configurations*

\mathbf{v} and \mathbf{v}' are considered neighboring if they differ in at most one user's contribution.

Definition 10 (LDP for L -neighboring). We say that a mechanism satisfies (ϵ, δ) -LDP for L -neighboring, iff it satisfies (ϵ, δ) -LDP w.r.t. the following L -neighboring notion: two input configurations \mathbf{v} and \mathbf{v}' are considered L -neighboring, iff the two vectors are otherwise identical except for at most one user's contribution v_i and v'_i ; and further, for the user i where the two vectors differ, it must be that $\|v_i - v'_i\|_1 \leq L$.

For the case of k -sparse binary vectors where each client's $v_i \in \{0, 1\}^d$, the following simple facts hold. A mechanism satisfies (ϵ, δ) -LDP for 1-neighboring, if and only if it is (ϵ, δ) -event-level-LDP. A mechanism satisfies (ϵ, δ) -LDP for k -neighboring, if and only if it satisfies (ϵ, δ) -user-level-LDP. More generally, for the case of k -sparse real-valued vectors where each client's $v_i \in [-1, 1]^d$, the following facts hold. If a mechanism satisfies (ϵ, δ) -LDP for 2-neighboring, it must also satisfy (ϵ, δ) -event-level-LDP. If a mechanism satisfies (ϵ, δ) -LDP for $2k$ -neighboring, it must also satisfy (ϵ, δ) -user-level-LDP.

Throughout the paper, unless otherwise noted, we use L_∞ -error to characterize the utility of our vector mean estimation mechanism. Specifically, L_∞ -error takes the maximum absolute error over all d coordinates.

For the special case where $k = 1$, Bassily and Smith [8] proved the following lower bound on the error of any (ϵ, δ) -event-level-LDP mechanism — note also that for the case $k = 1$, event-level and user-level LDP are the same up to a constant factor.

Theorem 11 (Lower bound on the error of single-item frequency estimation [8]). Suppose that $k = 1$. For any $\epsilon = O(1)$ and $0 \leq \delta \leq o(\frac{\epsilon}{n \log n})$, any non-interactive mechanism that satisfies (ϵ, δ) -event-level-LDP must incur expected L_∞ error of magnitude at least

$$\Omega\left(\min\left(\sqrt{\frac{\log(d)}{\epsilon^2 n}}, 1\right)\right).$$

IV. SPARSE VECTOR MEAN ESTIMATION

A. Algorithm

We give a unified algorithm that can be parametrized to achieve either event-level or user-level LDP, or LDP under L -neighboring. Our proposed algorithm is presented in Algorithm 1.

In the above algorithm, the clipping algorithm is needed only if we want to achieve $(\epsilon, 0)$ -DP under user-level LDP — see Section IV-D for more details. For all other cases, we achieve (ϵ, δ) -LDP.

Further, in the above algorithm, we assumed that the server computes the entire mean vector. However, when the universe size d is very large (e.g., the space of all possible URLs), the server may not want to write down the entire mean vector. Instead, it may wish to query \hat{v}_x for a specific item $x \in [d]$, e.g., the frequency of a specific URL. In this case, the server

Algorithm 1: Non-interactive Algorithm for k -Sparse Vector Mean Estimation

- 1 Parametrize bin number b , clipping range η , noise parameter Δ according to Table IV-A;
 - 2 **Client-side algorithm given input vector v_i :**
 - 3 Randomly pick a hash function $h_i : [d] \rightarrow [b]$
 - 4 Randomly pick a hash function $s_i : [d] \rightarrow \{-1, +1\}$
 - 5 **for** $j \in [b]$ **do**
 - 6 $B_{i,j} \leftarrow \sum_{l \in [d], h_i(l)=j} s_i(l)v_l$
 - 7 $\tilde{B}_{i,j} = \text{clip}_{[-\eta, +\eta]}(B_{i,j})$
 - 8 /*clipping needed only for pure user-level LDP */
 - 9 $\hat{B}_{i,j} \leftarrow \tilde{B}_{i,j} + \text{Lap}(\frac{\Delta}{\epsilon})$
 - 10 **end**
 - 11 Send $(h_i, s_i, \tilde{B}_{i,1}, \dots, \tilde{B}_{i,b})$ to the server
 - 12 **Server-side algorithm:**
 - 13 For all coordinate $x \in [d]$:
 $\hat{v}_x \leftarrow \frac{1}{n} \sum_{i \in [n]} s_i(x) \hat{B}_{i, h_i(x)}$
-

need not iterate through every $x \in [d]$, it only needs to invoke Line 13 for the items $x \in [d]$ that it cares about.

Discretizing real-numbers for transmission. In the above algorithm, we assumed that the client is transmitting real-valued numbers to the server. In practice, we can truncate and discretize real-valued numbers before transmitting, and ensure that the per-client communication cost is only $O(b \log k)$. The additional error introduced in the discretization process is asymptotically absorbed by the existing error terms, and therefore this step does not introduce any additional asymptotical error. See Appendix C-A for details.

Theorem 12 (Main theorem). Assuming $nk/b \geq \log(5d/\beta)$. Assuming the hash functions h and s are random oracles. Algorithm 1 satisfies L -neighboring (ϵ, δ) -LDP. With probability at least $1 - \beta$, the algorithm 1 outputs an estimation \hat{v} with L_∞ error of $O\left(\left(\sqrt{\frac{k}{b}} + \frac{\Delta}{\epsilon}\right) \sqrt{\frac{\log(d/\beta)}{n}}\right)$. The per-client communication cost is $O(b \log k)$.

Note that in practice, we can instantiate h and s with pseudorandom functions (PRFs) rather than random oracles. As mentioned earlier, the computational assumption here is not needed for the privacy but only for measure concentration.³

We present the privacy-related proof in Section IV-B and the utility-related proof in Section IV-C. For the communication cost analysis, see the discussion in Appendix C-A.

B. Privacy Analysis

Notice that the general L -neighboring LDP notion captures the requirement of event-level LDP($L = 2$) and user-level

³In theory, we can replace it with (almost) t -wise independent hash functions. For our measure concentration bounds to hold, we need t to be about logarithmically large, and a suitable hash function with a logarithmic description length, see [39]. We use PRF because it gives better constants for the concentration bound.

Cases	#Bins b	Clipping Range η	Laplacian Noise Magnitude Δ	L_∞ Error	Comm. Cost
Event-level LDP ($L = 2$)	$\frac{\epsilon^2 k}{4}$	∞	2	$O(\frac{1}{\epsilon} \sqrt{\frac{\log(d/\beta)}{n}})$	$O(k \log k)$
$L \leq \sqrt[3]{k}$	$\frac{\epsilon^2 k}{L^2}$	∞	L	$O(\frac{1}{\epsilon} L \sqrt{\frac{\log(d/\beta)}{n}})$	$O(\frac{k \log k}{L^2})$
$L \geq \sqrt[3]{k}$	$\sqrt{\frac{\epsilon^2 k}{L \log \frac{1}{\delta}}}$	∞	$3\sqrt{bL \log(2b/\delta)}$	$O(\frac{1}{\epsilon} (kL \log(k/L\delta))^{\frac{1}{4}} \sqrt{\frac{\log(d/\beta)}{n}})$	$O(\sqrt{\frac{k}{L}} \log k)$
User-level LDP ($L = 2k$)	1	$\sqrt{2k \log(4n/\beta)}$	2η	$O(\frac{1}{\epsilon} \sqrt{k \log(n/\beta)} \sqrt{\frac{\log(d/\beta)}{n}})$	$O(\log k)$

TABLE II

SUMMARY OF THE PARAMETERS UNDER DIFFERENT LDP SETTINGS, THE CORRESPONDING UTILITY GUARANTEE AND THE COMMUNICATION COSTS.

LDP($L = 2k$). Therefore, we only need to prove our algorithm is L -neighboring LDP and instantiate with corresponding L value for event- and user-level LDP. For now, we take the bin number b as an unspecified variable and we will provide the optimal selection of b later in the utility section.

Given two neighboring input configuration \mathbf{v}, \mathbf{v}' , from which one client's inputs are different, denoted as vectors v, v' . Then, $\|v - v'\|_1 \leq L$. We wish to bound the L_1 difference for the "raw bin values" B_1, \dots, B_b and B'_1, \dots, B'_b generated by two independent invocations of the client's algorithm.

Claim 13. *Given any two neighboring vectors v, v' . Taking the randomness of h and s , if $\Pr[\sum_{j \in [b]} |B_j - B'_j| \geq \Delta] \leq \delta$, then adding Laplacian noise of $\text{Lap}(\frac{\Delta}{\epsilon})$ ensures the two invocations of the client-side algorithm's output distributions are (ϵ, δ) -close.*

The proof is simple that one can compute the privacy budget loss in each bin and the total budget will be bounded by ϵ . We defer the proof to the appendix C-B.

For any h, s , rewrite $\sum_{j \in [b]} |B_j - B'_j| = \sum_{j \in [b]} \left| \sum_{l \in [d], h(l)=j} s(l)(v_j - v'_j) \right|$. With absolute inequality, the $s(l)$ term can be removed and the above expression is at most $\sum_{j \in [b]} \sum_{l \in [d], h(l)=j} |v_j - v'_j|$, which is exactly L . This proves the privacy property when $L \leq k^{\frac{1}{3}}$ where we set $\Delta = L$.

Moreover, we want to further prove that the difference after the binning is bounded by $\tilde{O}(\sqrt{bL})$. The intuition is that the binning process "squeezes" the difference, so that we can add smaller noise. For example, let's say a pair of neighboring vectors v, v' differ in coordinates x and y . Say $v_x = v_y = 1$ and $v'_x = v'_y = -1$. The original L_1 difference in the two coordinates are 4. Suppose the client samples a hash function h such that $h(x) = h(y)$. Then, we know that with probability $1/2$, the random ± 1 function s turns out to have outputs that $s(x) = -s(y)$. In this case, the influence in coordinates x and y , i.e., $|v_x - v'_x|$ and $|v_y - v'_y|$, cancel each other out in the bin because $|(s(x)v_x + s(y)v_y) - (s(x)v'_x + s(y)v'_y)| = 0$. We formally claim the following lemma:

Lemma 14. *Assuming $L/b \geq \log(2b/\delta)$. Consider any two neighboring vectors $v, v' \in [-1, +1]^d$ such that $\|v - v'\|_1 \leq L$. We have that $\Pr \left[\sum_{j \in [b]} |B_j - B'_j| > 3\sqrt{bL \log(2b/\delta)} \right] \leq \delta$.*

Proof. Fix a bin j . Define random variables $Z_l = \mathbb{I}[h(l) = j]s(l)(v_l - v'_l)$ for $l \in [d]$ and $Z = \sum_{l \in [d]} Z_l$. Z 's distribution

is exactly the difference in bin j after binning. We know that the variables $\{Z_l\}_{l \in [d]}$ are independent and bounded by $[-2, 2]$. Also, $\mathbf{E}[Z_l^2] = \frac{1}{b}(v_l - v'_l)^2 \leq \frac{2}{b}(v_l - v'_l)$. Let $\mu = \frac{L}{b}$. Using Bernstein's inequality, setting $t = 3\sqrt{\mu \log(2b/\delta)}$, we have

$$\begin{aligned} \Pr[|Z| \geq t] &\leq 2 \exp \left(- \frac{\frac{1}{2}t^2}{\sum_{l \in [d]} \mathbf{E}[Z_l^2] + \frac{2}{3}t} \right) \\ &\leq 2 \exp \left(- \frac{\frac{9}{2}\mu \log(2b/\delta)}{2\mu + 2\sqrt{\mu \log(2b/\delta)}} \right) = 2(\delta/2b)^{\frac{\frac{9}{2}\mu}{2\mu + 2\sqrt{\mu \log(2b/\delta)}}} \end{aligned}$$

Using the condition that $\mu = \frac{L}{b} \geq \log(b/\delta)$, we have $\frac{\frac{9}{2}\mu}{2\mu + 2\sqrt{\mu \log(2b/\delta)}} \geq \frac{9}{8} \geq 1$. Therefore, $\Pr \left[|Z| \geq 3\sqrt{\mu \log(2b/\delta)} \right] \leq \delta/b$. Taking the union bound over all b bins, we have the total difference in all bins are at most $3\sqrt{bL \log(2b/\delta)}$ with probability at least $1 - \delta$. \square

By Claim 13 and Lemma 14, setting the noise parameter to $\Delta = 3\sqrt{bL \log(2b/\delta)}$ is enough to achieve (ϵ, δ) -privacy under L -Neighboring LDP, assuming $L/b \geq \log(2b/\delta)$. In practice, when we search for the optimal b , we will carefully set b such that the condition $L/b \geq \log(2b/\delta)$ is held.

C. Utility Analysis

We first provide the simplified version of the utility part for the main theorem for general parameter settings – bin number b , clipping range η and the noise parameter Δ . The full proof is deferred to Appendix C-C. Then, we will discuss how to choose the optimal b to achieve the best utility under different scenarios.

Proof. (Sketch) Fix an index $x \in [d]$. The server computes the estimation \hat{v}_x as $\hat{v}_x = \frac{1}{n} \sum_{i \in [n]} s_i(x) \tilde{B}_{i, h_i(x)}$. We bound the error by the three steps: 1) binning; 2) clipping; 3) adding Laplacian noise.

Binning error. The absolute error of binning is $\left| \frac{1}{n} \sum_{i \in [n]} v_{i,x} - \frac{1}{n} \sum_{i \in [n]} B_{i, h_i(x)} s_i(x) \right|$. Define random variables $Y_{i,l} = \mathbb{I}[h_i(l) = h_i(x)] v_{i,l} s_i(x)$ as the error introduced by coordinate $l \neq x$ in client i 's vector. The error is equal to $\left| \frac{1}{n} \sum_{i \in [n]} \sum_{l \in [d], l \neq x} Y_{i,l} \right|$. Since we model the hash function h_i as random oracle, the hash collision probability is $\frac{1}{b}$. Also, the hash function s_i is a uniform ± 1 function, so we have $Y_{i,l} = v_{i,l}$ with prob. $\frac{1}{2b}$ and $Y_{i,l} = -v_{i,l}$ with prob. $\frac{1}{2b}$. We consider all nk

non-zero coordinates and we can use the analysis for a zero-mean random walk with total length of $\frac{nk}{b}$. Using Bernstein’s inequality, we can prove that for all coordinate x , $|\frac{1}{n} \sum_{i \in [n]} v_{i,x} - \frac{1}{n} \sum_{i \in [n]} B_{i,h_i(x)} s_i(x)| = O(\sqrt{\frac{k}{b}} \sqrt{\frac{\log(d/\beta)}{n}})$ with probability at least $1 - O(\beta)$.

Clipping error. We actually prove that the clipping range is large enough, so that the clipping error is zero with high probability. We directly compute the raw bin value $B_{i,j}$ ’s moment generating function and conclude that it is a sub-Gaussian r.v. with a variance at most k . That means the absolute bin values will roughly be $\tilde{O}(\sqrt{k})$. Set the clipping range to $\eta \geq \sqrt{2k \log(4nb/\beta)}$. Using concentration bound for sub-Gaussian variables and taking union bound over all nb bins across n clients, we conclude the probability of the clipping error being zero is at least $1 - O(\beta)$.

Laplacian noise error. Finally, we look at the absolute error term introduced by adding Laplacian noise. It turns out that the error’s distribution is the same as the distribution for the mean of n i.i.d. Laplacian variables with parameter $\frac{\Delta}{\epsilon}$. Using the concentration bound for Laplacian noise and taking the union bound over all $x \in [d]$, the maximal error is bounded by $O\left(\frac{\Delta}{\epsilon} \sqrt{\frac{\log(d/\beta)}{n}}\right)$ with probability $1 - O(\beta)$.

Combine the above arguments. By taking the union bound and setting the constants appropriately, we can conclude that the L_∞ error is $O\left(\left(\sqrt{\frac{k}{b}} + \frac{\Delta}{\epsilon}\right) \sqrt{\frac{\log(d/\beta)}{n}}\right)$ with probability $1 - \beta$. \square

Selecting the number of bins b . We show how to set the number of bins b to minimize error. Fix k and L . From the privacy analysis section, we know that Δ can be set to $\min(L, 3\sqrt{bL \log(2b/\delta)})$. Define functions $f_1(b) = \sqrt{\frac{k}{b}} + \frac{L}{\epsilon}$ and $f_2(b) = \sqrt{\frac{k}{b}} + \sqrt{bL \log(2b/\delta)}/\epsilon$. The error can be rewritten as $\min(f_1(b) + f_2(b)) \cdot \tilde{O}(\frac{1}{\sqrt{n}})$. Optimizing $f_1(b)$, the optimal b is $b_1^* = \frac{\epsilon^2 k}{L^2}$ and $f_1(b_1^*) = \frac{L}{\epsilon}$. Optimizing $f_2(b)$, the optimal b is roughly $b_2^* = \sqrt{\frac{\epsilon^2 k}{L \log \frac{1}{\delta}}}$ and $f_2(b_2^*) = O(\frac{1}{\epsilon}(kL \log(\frac{k}{L\delta}))^{\frac{1}{4}})$. Then, by comparing two local minima, we conclude that when $L \leq k^{\frac{1}{3}}$, the optimal b is $\frac{\epsilon^2 k}{L^2}$ and the error is $O\left(\frac{L}{\epsilon} \sqrt{\frac{\log(d/\beta)}{n}}\right)$. When $L \geq k^{\frac{1}{3}}$, the optimal b is roughly $\sqrt{\frac{\epsilon^2 k}{L \log \frac{1}{\delta}}}$ and the error is $O\left(\frac{1}{\epsilon}(kL \log(\frac{k}{L\delta}))^{\frac{1}{4}} \sqrt{\frac{\log(d/\beta)}{n}}\right)$. In practice, we also take concrete constants into consideration that we perform a parameter search for b , predict the final error and select the best b accordingly.

D. Achieving $(\epsilon, 0)$ -user-level LDP

We analyze how the extra clipping step in Algorithm 1 achieves pure-LDP in user-level setting. The idea is to push the “failure probability” δ in privacy definition to the utility

Datasets	#Clients n	#Items d	#Records	Sparsity k
Clothing [44]	47958	1378	79285	6
Renting [42]	105571	5850	183052	11
Movies [45]	138493	26744	7019990	100

TABLE III
REAL-WORLD DATASET.

theorem’s failure probability β . We first observe that in user-level LDP, we have the neighboring distance $L = 2k$ and the optimal bin number selection is $b = 1$. We now consider the magnitude of the “raw bucket value” $B_{i,1}$ for client each i . We simply have $B_{i,1} = \sum_{l \in [d]} v_{i,l} s_i(l)$. Take the randomness of the random ± 1 function s_i . We can see the distribution of $B_{i,1}$ is similar to a zero-mean random walking with at most k steps, where each step’s length is at most 1. Using Bernstein-type concentration bound, we can prove that with probability $1 - O(\beta)$, for all client $i \in [n]$, $|B_{i,1}| \leq \sqrt{2k \log(4n/\beta)}$. See the detail proof in Appendix C-C. Let $\eta = \sqrt{2k \log(4n/\beta)}$. We see that the difference in $B_{i,1}$ and $B'_{i,1}$ in two independent invocation of the client-side algorithm given input v, v' are at most 2η with probability 1 after clipping. That means we only need to set the noise parameter $\Delta = 2\eta$ and the algorithm is $(\epsilon, 0)$ -user-level-LDP. Plug the parameters into the main theorem, we know the utility guarantee of this optimization is $O(\frac{1}{\epsilon} \sqrt{k \log(n/\beta)} \sqrt{\frac{\log(d/\beta)}{n}})$. We see the utility guarantee of this optimization is similar to the original unclipped version – they are both $\tilde{O}(\frac{1}{\epsilon} \frac{\sqrt{k}}{\sqrt{n}})$. In practice, this clipped version has much smaller constant factor in terms of error.

V. EVALUATION

A. Setup

Implementation. To evaluate our approach, we implement it with C++, compile it with gcc4.8 and the C++11 standard. We use 40-bit random seeds to generate the hash functions. For simplicity, we directly use 32-bit floating numbers to store and transmit real values.

Datasets. We evaluate the algorithms for both synthetic and real-world datasets. For the synthetic dataset, we assume there are 10^5 users, each with a vector of dimension d and sparsity k . We first randomly sample the non-zero coordinates according to Zipf’s distribution with a suitable degrading parameter ($s = 1.4$). We choose the Zipf’s distribution because it naturally appears in real-world data analytics. For each sampled non-zero coordinate, the actual value is sampled from a Gaussian distribution with mean $\mu = 1$ and standard deviation $\sigma = 0.3$. Then the values are clipped to $[-1, 1]$.

For the real-world dataset experiment, we downloaded three open-sourced datasets from Kaggle, including an online cloth shopping dataset [44], a clothing renting dataset [42] and a movie rating dataset [45], where each record describes one activity (purchase, rent, or rating, respectively). Table V-A gives more information about the datasets. We select those records with client feedback ratings and normalize them to $[-1, 1]$. Given the sparsity parameter is k , for clients with more than k records, we randomly sample k records.

Metrics. We consider both utility and communication cost fixing the privacy level (i.e., fixing ϵ and δ). To measure utility, we use the L_∞ error and the mean square error (MSE). Given $\hat{v} = \frac{1}{n} \sum_{i \in [n]} v_i$ as the true mean vector and \bar{v} as the estimation vector, they are defined as:

$$L_\infty \text{ Error} = \max_{x \in [d]} |\hat{v}_x - \bar{v}_x| \quad \text{MSE} = \frac{1}{d} \sum_{x \in [d]} (\hat{v}_x - \bar{v}_x)^2$$

For the communication cost, we measure the per-client communication cost: We sum up the **byte-length** of all the reports from the clients and compute the average report size.

Evaluation Roadmap. We split the experiments into three groups: user-level LDP setting, event-level setting LDP, and the L -Neighboring setting. Within each group, we measure different methods varying three parameters: sparsity k , privacy budget ϵ (in most cases, we use $\delta = 0$; but when $\delta > 0$, e.g., for the naive perturbation scheme with Gaussian noise, we always use $\delta = 10^{-5}$), and dimension size d . We mainly compare our proposed method with the k -fold repetition-plus-1-sparse mechanism (referred as *k-fold repetition*), the sampling + 1-sparse mechanism (referred as *sampling*), the naive perturbation mechanism (with Gaussian Noise [6]), Harmony [43] and PCKV [30]. We run the experiment 10 times and report the average error and the average communication cost.

B. Performance under User-level LDP

User-level LDP is the more standard setting in LDP analytics. Existing methods are mostly designed for user-level LDP. We first compare our method against existing ones in this setting.

Varying sparsity k . We plot the L_∞ error results in Figure 1(a) and the MSE results Figure 1(d). In our theoretical analysis, we prove that the L_∞ error of our algorithm scales with \sqrt{k} . The sampling + 1-sparse method's error scales with k , and other algorithm cannot utilize the sparsity. The figures show that our method has the smallest estimation error for the whole region when k ranges from 1 to 1024. The error of the sampling solution and the naive perturbation mechanism scales with k and they perform worse than PCKV and Harmony when the sparsity k is larger than \sqrt{d} .

Varying privacy budget ϵ . The results are shown in Figure 1(b) and Figure 1(e). With larger privacy budget, all schemes except Harmony achieve better estimation errors. However, when the dimension d is sufficiently large, Harmony and PCKV suffer from a $\tilde{O}(\frac{\sqrt{d}}{\sqrt{n}})$ error. In the relatively high privacy budget region, PCKV shows better performance. Our method always has the smallest error in the reasonable large privacy budget range.

Varying dimension d . In many use cases, the domain size (vector length) can be extremely huge, such as all possible products on Amazon, all possible URL and all geographical location on the earth. In this experiment, we only measure the top 100 coordinate with the largest absolute mean value. This

is actually inspired by a real use case where the domain size is sufficiently and the server only wishes to compute the value for a limited keys (e.g. website access analysis). The results are shown in Figure 1(b) and Figure 1(e). Our method provides an important feature – its utility and communication cost decouple from the domain size. Our method can maintain a stable estimation error even with very large dimension d , while using minimum communication cost. The naive perturbation scheme needs to communicate $O(d)$ bits between the clients and the server. In the very dense case, where $k \approx d$, PCKV and Harmony has slightly better estimation error because our method has the extra $\sqrt{\log n}$ term in the error. However, in the more sparse case, all other methods fail to provide any meaningful guess. The noticeable drop in the large d region of the error curves for PCKV and Harmony is because they basically output a meaningless zero vector.

C. Performance under Event-level LDP

The results are plotted in Figure 2. Theoretically (from Table I), our method is better than other methods by at least a polynomial gap \sqrt{k} in terms of the L_∞ error. The following experiments verify the theoretical results.


Varying sparsity k . The results are shown in Figure 2(a) and Figure 2(d). The results matches our theoretical results that our methods are not scale with k in event-level LDP. It also outperforms other methods for the whole range.

Varying privacy budget ϵ . The results are shown in Figure 2(b) and Figure 2(e). The higher privacy budget are beneficial to all methods' utility performance. Our algorithm still has the best estimation performance for the whole range.

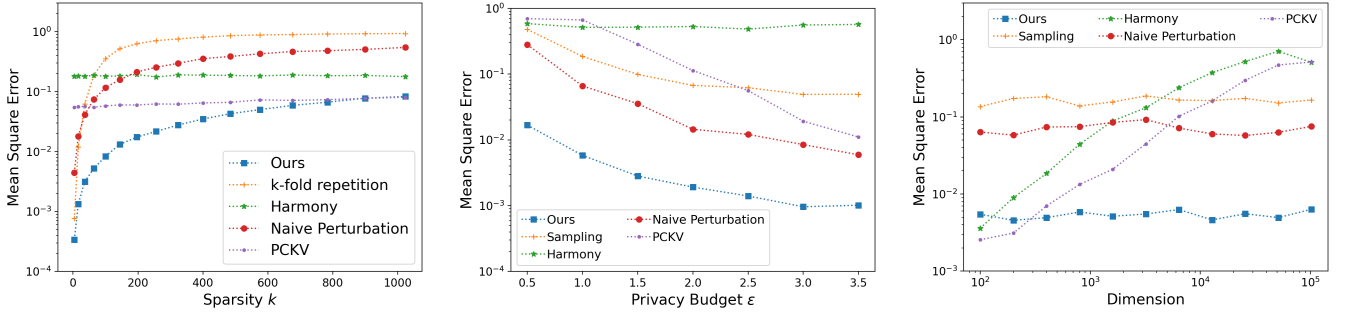
Varying dimension d . The results are shown in Figure 2(c) and Figure 2(f). Again, our algorithm has decoupled from the dimension d and it has much smaller error estimation than other methods.

D. Performance under L-Neighboring LDP

The neighboring distance L provides a better way to describe the middle ground between user-level LDP and event-level LDP. Our algorithm has theoretical L_∞ error of $\min\{O(L), O((kL \log(1/\delta))^{\frac{1}{4}}), O(\sqrt{k \log n})\} \cdot O(\frac{1}{\epsilon} \sqrt{\frac{\log d}{n}})$. In the experiment, we fix the sparsity $k = 64$ and vary the neighboring L_1 distance from 1 to 128. In the case when $L \ll k$, the parameter configuration with $O((kL \log(1/\delta))^{\frac{1}{4}})$ error growing factor should have asymptotically advantage over the configuration with $O(\sqrt{k \log n})$ growing factor. However, in practice, we realize that the latter scheme(the algorithm with clipping) has a much smaller constant factor. Hence, in the case when k is not large enough, we only see the optimized clipping scheme dominates the unclipped scheme. The mixed strawman solutions, including k -fold repetition scheme and sampling scheme, can only adapt to either event-level LDP or user-level LDP. PCKV and Harmony cannot fully utilize the relaxed privacy as a way to improve the estimation error. The naive perturbation mechanism has worse scaling factor than



(a) L_∞ error results when fixing $n = 10^5$, $d = 4096$, $\epsilon = 1.0$ and varying k from 1 to 1024. (b) L_∞ error results for top 100 coordinates when fixing $n = 10^5$, $d = 10^5$, $k = 64$ and varying ϵ from 0.5 to 3.5. (c) L_∞ error results for top 100 coordinates when fixing $n = 10^5$, $k = 64$, $\epsilon = 1.0$ and varying d from 64 to 10^5 .



(d) MSE results when fixing $n = 10^5$, $d = 4096$, $\epsilon = 1.0$ and varying k from 1 to 1024. (e) MSE results results for top 100 coordinates when fixing $n = 10^5$, $d = 10^5$, $k = 64$ and varying ϵ from 0.5 to 3.5. (f) MSE results for top 100 coordinates when fixing $n = 10^5$, $k = 64$, $\epsilon = 1.0$ and varying d from 64 to 10^5 .

Fig. 1. Comparing the utilities of our method and existing approaches under user-level LDP. The subfigures in the top and bottom row show the results of L_∞ error and MSE, respectively. The subfigures in the left, middle, and right column varies sparsity k (from 1 to 1024), ϵ (from 0.5 to 3.5), and dimension d (from 64 to 10^5), respectively (while fixing the other two parameters).

Name	Event-level LDP			User-level LDP		
	Comm. Cost	L_∞ Err.	MSE	Comm. Cost	L_∞ Err.	MSE
k -fold repetition	28	0.092	0.00059	-	-	-
Sampling	-	-	-	8	0.18	0.0034
Naive Perturbation	5512	0.17	0.0023	5512	0.41	0.014
Harmony	-	-	-	8	1.0	0.13
PCKV	-	-	-	8	0.70	0.038
Ours	9	0.040	0.00019	9	0.098	0.00095

TABLE IV
EXPERIMENT RESULTS FOR THE CLOTHING DATASET.

Name	Event-level LDP			User-level LDP		
	Comm. Cost	L_∞ Err.	MSE	Comm. Cost	L_∞ Err.	MSE
k -fold repetition	48	0.085	0.00050	-	-	-
Sampling	-	-	-	8	0.29	0.0052
Naive Perturbation	23400	0.12	0.0010	23400	0.40	0.012
Harmony	-	-	-	8	1.0	0.24
PCKV	-	-	-	8	1.0	0.074
Ours	13	0.033	0.00083	9	0.11	0.00091

TABLE V
EXPERIMENT RESULTS FOR THE RENTING DATASET.

Name	Event-level LDP			User-level LDP		
	Comm. Cost	L_∞ Err.	MSE	Comm. Cost	L_∞ Err.	MSE
k -fold repetition	404	0.25	0.0033	-	-	-
Sampling	-	-	-	8	1.0	0.29
Naive Perturbation	109052	0.12	0.00082	109052	1.0	0.081
Harmony	-	-	-	8	1.0	0.49
PCKV	-	-	-	8	1.0	0.24
Ours	105	0.034	0.000065	9	0.25	0.0021

TABLE VI
EXPERIMENT RESULTS FOR THE MOVIE DATASET.

the L_∞ error, which is roughly 2.4, 3.3 and 8.0 in the three datasets correspondingly. The Harmony and PCKV schemes do not output very meaningful estimation in the experiments because their algorithms have error scaled with the dimension d .

VI. LOWER BOUND


In a previous work [8], Bassily and Smith showed a lower bound of $\Omega\left(\frac{1}{\epsilon} \sqrt{\frac{\log d}{n}}\right)$ on the L_∞ error under the 1-sparse case with the constraints of $(\epsilon, o(\frac{1}{n \log n}))$ -LDP (Theorem 11). The 1-sparse case can be seen as a special case for the general k -sparse vector mean estimation under event-level LDP. Our algorithm for event-level LDP matches this lower bound, making the error bound tight in the event-level LDP case.

We observe that it is not hard to extend the framework and prove a lower bound of $\Omega\left(\frac{1}{\epsilon} \sqrt{\frac{k \log(d/k)}{n}}\right)$ on the L_∞ error

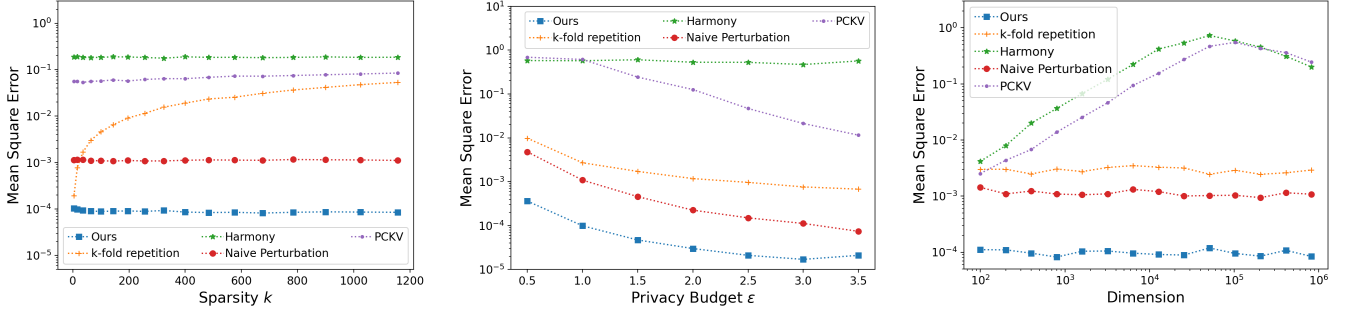
our method, but in the turning point where $L = \sqrt{k \log n}$, it roughly matches the error of our method.

E. Real-world Dataset Experiments

We compile the Clothing, Renting and Movie dataset to the sparse vector mean estimation problem. The description of the datasets can be found in Table V-E. Our method achieves best accuracy in both event-level LDP setting and user-level LDP setting by a magnitude of gap. Specifically, compared to our method, the strawman scheme has an extra \sqrt{k} factor in



(a) L_∞ error results when fixing $n = 10^5$, $d = 4096$, $\epsilon = 1.0$ and varying k from 1 to 1024. (b) L_∞ error results for top 100 coordinates when fixing $n = 10^5$, $d = 10^5$, $k = 64$ and varying ϵ from 0.5 to 3.5. (c) L_∞ error results for top 100 coordinates when fixing $n = 10^5$, $k = 64$, $\epsilon = 1.0$ and varying d from 64 to 10^5 .



(d) MSE results when fixing $n = 10^5$, $d = 4096$, $\epsilon = 1.0$ and varying k from 1 to 1024. (e) MSE results for top 100 coordinates when fixing $n = 10^5$, $d = 10^5$, $k = 64$ and varying ϵ from 0.5 to 3.5. (f) MSE results for top 100 coordinates when fixing $n = 10^5$, $k = 64$, $\epsilon = 1.0$ and varying d from 64 to 10^5 .

Fig. 2. Comparing the utilities of our method and existing approaches under event-level LDP. The plotting convention follows that of Figure 1: subfigures in the top and bottom row show the results of L_∞ error and MSE, respectively. The subfigures in the left, middle, and right column varies sparsity k (from 1 to 1024), ϵ (from 0.5 to 3.5), and dimension d (from 64 to 10^5), respectively (while fixing the other two parameters).

of k -sparse vector mean estimation under the user-level LDP. For completeness, we present the full proof below.

Notation. In the lower bound proof, each client i has a k -sparse input vector $v_i \in \mathcal{S} := \{v \in \{0, 1\}^d : \|v\|_1 = k\}$, where the special case $k = 1$ is essentially the one-item frequency estimation problem [8]. Note that since this setting is a special case of real-valued mean vector estimation, the lower bound applies to mean vector estimation more generally.

Each client i applies an (ϵ, δ) -differentially private (where any two inputs in \mathcal{S} are neighboring) algorithm $\mathcal{Q}_i(\cdot)$ independently to produce $\mathbf{z}_i = \mathcal{Q}_i(v_i)$ in some report space \mathcal{Z} . The server computes $\hat{v} := \mathcal{A}(\mathbf{z}_1, \dots, \mathbf{z}_n)$, which estimates $\frac{1}{n} \sum_i v_i$. Then the following lower bound holds.

Theorem 15 (Lower bound on error, k -sparse mean vector estimation). *Let $0 < \epsilon = O(1)$ and $0 < \delta = o(\frac{\epsilon}{n \log n})$. Suppose for each client i , the (randomized) algorithm $\mathcal{Q}_i : \mathcal{S} \rightarrow \mathcal{Z}$ is (ϵ, δ) -differentially private, where any two inputs in \mathcal{S} are considered as neighboring. Moreover, $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathcal{S}$ is a (potentially randomized) aggregator function.*

Then, there exists some distribution \mathcal{P} on \mathcal{S} (depending on \mathcal{Q}_i 's and \mathcal{A}) such that if every client i independently generates a report $\mathbf{z}_i = \mathcal{Q}_i(v_i)$, where v_i is sampled from \mathcal{P} independently, the expected error of estimating $\bar{v} := \frac{1}{n} \sum_i v_i$ has the following lower bound:

$$\mathbb{E}[\|\mathcal{A}(\mathbf{z}_1, \dots, \mathbf{z}_n) - \bar{v}\|_\infty] \geq \min \left\{ \Omega \left(\frac{1}{\epsilon} \sqrt{\frac{\log |\mathcal{S}|}{n}} \right), 1 \right\},$$

where $\log |\mathcal{S}| = \log \binom{d}{k} \geq k \log(d/k)$.

Plugging in the definition of \mathcal{S} , the following corollary gives our main lower bound for user-level-LDP.

Corollary 16. *Observing that if $L = 2k \leq \sqrt{d}$, then any two inputs in \mathcal{S} has L_∞ distance at most L . Hence, in this case,*

$$\mathbb{E}[\|\mathcal{A}(\mathbf{z}_1, \dots, \mathbf{z}_n) - \bar{v}\|_\infty] \geq \min \left\{ \Omega \left(\frac{1}{\epsilon} \sqrt{\frac{L \log d}{n}} \right), 1 \right\}$$

Proof roadmap. Just like Bassily and Smith, our goal is to find a “hard” joint distribution \mathcal{P} on the clients’ inputs $\mathbf{v} = (v_1, \dots, v_n)$, such that the expected L_∞ estimation error is large for any (ϵ, δ) -user-level LDP algorithm. Here, the expectation is taken over the randomness coming from the input sampling and the algorithm. We construct the distribution \mathcal{P} as following. First, a vector $V \in \{0, 1\}^d$ is sampled uniformly at random from a candidate set \mathcal{S} that includes all binary k -sparse vectors in $\{0, 1\}^d$. Next, each client’s input v_i is sampled i.i.d from a distribution $\mathcal{P}_V^{(\eta)}$ (using the same V for all users) as follows:

$$v_i = \begin{cases} V & \text{w.p. } \eta \\ U & \text{w.p. } 1 - \eta \end{cases} \quad (2)$$

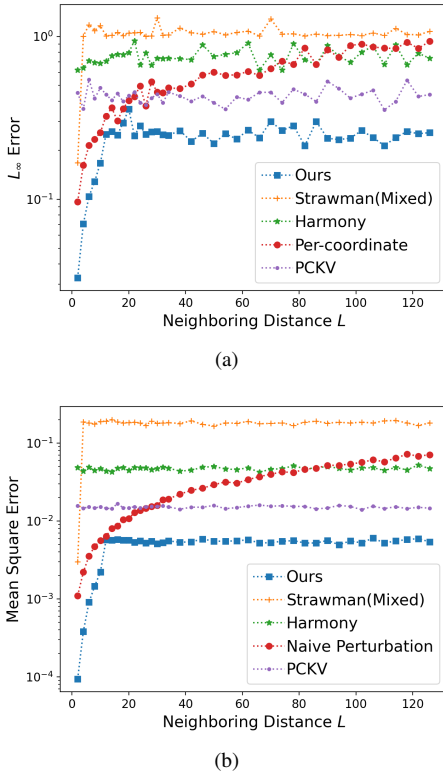


Fig. 3. Comparing the utilities of our method and existing approaches under L -Neighboring LDP.

where U is drawn uniformly from \mathcal{S} . The distribution $\mathcal{P}_V^{(\eta)}$ is an instance of an η -degrading channel [8]. To prove that any (ϵ, δ) -LDP algorithm has large error with respect to $\mathcal{P}_v^{(\eta)}$ for at least one $v \in \mathcal{S}$, we view the problem as an encoding-decoding process, then bound the error using Fano's inequality. Each client i generates a report $\mathbf{z}_i = Q_i(v_i)$. The joint reports $\mathbf{z} := (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n)$ are viewed as a noisy encoding of V . In an attempt to recover V , the server aggregator function \mathcal{A} is applied to produce the mean estimation $\mathcal{A}(\mathbf{z})$. Then, to decode the original V , the server removes the bias introduced by the degrading channel then rounds the estimation $\mathcal{A}(\mathbf{z})$ to the nearest binary vector \hat{V} . A decoding error occurs if $\hat{V} \neq V$.

The lower bound proof relies on two bounds on the probability of decoding error. On one hand, the differential privacy of each Q_i implies that the mutual information $I(V; \mathbf{z})$ is small, which means that the decoding error probability is large by Fano's inequality. On the other hand, a small L_∞ -error estimation of the mean vector implies that the original V can be recovered from $\mathcal{A}(\mathbf{z})$ with high probability. These effects limit the decoding error probability and give us a lower bound on the mean vector estimation error. That is, for small enough $\delta = o(\frac{\epsilon}{n \log n})$, we find a distribution \mathcal{P} over candidate set \mathcal{S} that implies a lower bound of $\Omega(\frac{1}{\epsilon} \sqrt{\frac{\log |\mathcal{S}|}{n}})$ on the L_∞ -error of mean vector estimation. By considering k -sparse vectors in $\{0, 1\}^d$, (for which $|\mathcal{S}| = \binom{d}{k}$), we obtain the lower bound of $\Omega\left(\frac{1}{\epsilon} \sqrt{\frac{k \log(d/k)}{n}}\right)$.

In the interest of space, we defer the detailed lower bound proof to Appendix D.

ACKNOWLEDGMENTS

This work is supported in part by NSF awards under the grant numbers 2128519, 2044679, CIF-1705007, and CA-2040675, as well as a grant from ONR and AFOSR grant FA9550-21-1-0090. The authors also acknowledge support from a Packard Fellowship, the Sloan Foundation, and gift grants from Cisco, Siemens, Google, Bosch, and J.P. Morgan Chase. T-H. Hubert Chan was partially funded by the Hong Kong RGC under the grants 17200418 and 17201220.

REFERENCES

- [1] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [2] J. Acharya, Z. Sun, and H. Zhang, "Hadamard response: Estimating distributions privately, efficiently, and with little communication," *AISTATS*, 2018.
- [3] J. Acharya and Z. Sun, "Communication complexity in locally private distribution estimation and heavy hitters," *CoRR*, vol. abs/1905.11888, 2019.
- [4] M. Aumüller, C. J. Lebeda, and R. Pagh, "Differentially private sparse vectors with low error, optimal space, and fast access," *arXiv preprint arXiv:2106.10068*, 2021.
- [5] E. Bagdasaryan, P. Kairouz, S. Mellem, A. Gascón, K. Bonawitz, D. Estrin, and M. Gruteser, "Towards sparse federated analytics: Location heatmaps under distributed differential privacy with secure aggregation," *arXiv preprint arXiv:2111.02356*, 2021.
- [6] B. Balle and Y.-X. Wang, "Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising," in *International Conference on Machine Learning*. PMLR, 2018, pp. 394–403.
- [7] R. Bassily, K. Nissim, U. Stemmer, and A. Thakurta, "Practical locally private heavy hitters," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17, 2017, p. 2285–2293.
- [8] R. Bassily and A. Smith, "Local, private, efficient protocols for succinct histograms," in *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, 2015, pp. 127–135.
- [9] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers, "Protection against reconstruction and its applications in private federated learning," 2019.
- [10] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 1175–1191.
- [11] D. Boneh, E. Boyle, H. Corrigan-Gibbs, N. Gilboa, and Y. Ishai, "Lightweight techniques for private heavy hitters," in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 762–776.
- [12] M. Bun, J. Nelson, and U. Stemmer, "Heavy hitters and the structure of local privacy," in *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Houston, TX, USA, June 10-15, 2018*. ACM, 2018, pp. 435–447.
- [13] X. Cao, J. Jia, and N. Z. Gong, "Data poisoning attacks to local differential privacy protocols," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 947–964.
- [14] T.-H. H. Chan, E. Shi, and D. Song, "Private and continual release of statistics," *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 3, pp. 1–24, 2011.
- [15] —, "Privacy-preserving stream aggregation with fault tolerance," in *Financial Cryptography and Data Security (FC)*, 2012.
- [16] M. Charikar, K. Chen, and M. Farach-Colton, "Finding frequent items in data streams," in *International Colloquium on Automata, Languages, and Programming*. Springer, 2002, pp. 693–703.
- [17] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate, "Differentially private empirical risk minimization," *J. Mach. Learn. Res.*, vol. 12, jul 2011.

- [18] W. Chen, P. Kairouz, and A. Özgür, “Breaking the communication-privacy-accuracy trilemma,” in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [19] G. Cormode, T. Kulkarni, and D. Srivastava, “Marginal release under local differential privacy,” in *Proceedings of the 2018 International Conference on Management of Data*, 2018, pp. 131–146.
- [20] G. Cormode, C. Procopiuc, D. Srivastava, and T. T. Tran, “Differentially private summaries for sparse data,” in *Proceedings of the 15th International Conference on Database Theory*, 2012, pp. 299–311.
- [21] H. Corrigan-Gibbs and D. Boneh, “Prio: Private, robust, and scalable computation of aggregate statistics,” in *14th USENIX symposium on networked systems design and implementation (NSDI 17)*, 2017, pp. 259–282.
- [22] B. Ding, J. Kulkarni, and S. Yekhanin, “Collecting telemetry data privately,” *arXiv preprint arXiv:1712.01524*, 2017.
- [23] J. Duchi and R. Rogers, “Lower bounds for locally private estimation via communication complexity,” in *Proceedings of the Thirty-Second Conference on Learning Theory*, vol. 99, 2019, pp. 1161–1191.
- [24] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, “Minimax optimal procedures for locally private estimation,” *Journal of the American Statistical Association*, vol. 113, no. 521, pp. 182–201, 2018.
- [25] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Theory of cryptography conference*. Springer, 2006, pp. 265–284.
- [26] C. Dwork and A. Roth, “The algorithmic foundations of differential privacy,” *Foundations and Trends in Theoretical Computer Science*, vol. 9, no. 3-4, pp. 211–407, 2014. [Online]. Available: <http://dblp.uni-trier.de/db/journals/ftcs/ftcs9.html#DworkR14>
- [27] 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.
- [28] G. Fanti, V. Pihur, and Ú. Erlingsson, “Building a rappor with the unknown: Privacy-preserving learning of associations and data dictionaries,” *arXiv preprint arXiv:1503.01214*, 2015.
- [29] V. Feldman, I. Mironov, K. Talwar, and A. Thakurta, “Privacy amplification by iteration,” in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2018, pp. 521–532.
- [30] X. Gu, M. Li, Y. Cheng, L. Xiong, and Y. Cao, “Pckv: Locally differentially private correlated key-value data collection with optimized utility,” in *29th USENIX Security Symposium*, 2020, pp. 967–984.
- [31] M. Hardt and G. N. Rothblum, “A multiplicative weights mechanism for privacy-preserving data analysis,” in *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE, 2010, pp. 61–70.
- [32] P. Kairouz, Z. Liu, and T. Steinke, “The distributed discrete gaussian mechanism for federated learning with secure aggregation,” *arXiv preprint arXiv:2102.06387*, 2021.
- [33] P. Kairouz, S. Oh, and P. Viswanath, “The composition theorem for differential privacy,” in *International conference on machine learning*. PMLR, 2015, pp. 1376–1385.
- [34] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith, “What can we learn privately?” *SIAM Journal on Computing*, vol. 40, no. 3, pp. 793–826, 2011.
- [35] J. Konecny, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [36] A. Korolova, K. Kenthapadi, N. Mishra, and A. Ntoulas, “Releasing search queries and clicks privately,” in *Proceedings of the 18th international conference on World wide web*, 2009, pp. 171–180.
- [37] Z. Li, T. Wang, M. Lopuhaä-Zwakenberg, N. Li, and B. Skoric, “Estimating numerical distributions under local differential privacy,” in *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 621–635. [Online]. Available: <https://doi.org/10.1145/3318464.3389700>
- [38] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [39] R. Meka, O. Reingold, G. N. Rothblum, and R. D. Rothblum, “Fast pseudorandomness for independence and load balancing,” in *International Colloquium on Automata, Languages, and Programming*. Springer, 2014, pp. 859–870.
- [40] I. Mironov, “Rényi differential privacy,” in *2017 IEEE 30th computer security foundations symposium (CSF)*. IEEE, 2017, pp. 263–275.
- [41] I. Mironov, O. Pandey, O. Reingold, and S. P. Vadhan, “Computational differential privacy,” in *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, ser. Lecture Notes in Computer Science, S. Halevi, Ed., vol. 5677. Springer, 2009, pp. 126–142.
- [42] R. Misra, “Clothing fit dataset for size recommendation,” Aug 2018. [Online]. Available: <https://www.kaggle.com/rmisra/clothing-fit-dataset-for-size-recommendation>
- [43] T. T. Nguyễn, X. Xiao, Y. Yang, S. C. Hui, H. Shin, and J. Shin, “Collecting and analyzing data from smart device users with local differential privacy,” *arXiv preprint arXiv:1606.05053*, 2016.
- [44] Nicapotato, “Women’s e-commerce clothing reviews,” Feb 2018. [Online]. Available: <https://www.kaggle.com/nicapotato/womens-e-commerce-clothing-reviews>
- [45] Pooh, “Movie rating data,” 2017. [Online]. Available: <https://www.kaggle.com/ashukr/movie-rating-data>
- [46] Z. Qin, Y. Yang, T. Yu, I. Khalil, X. Xiao, and K. Ren, “Heavy hitter estimation over set-valued data with local differential privacy,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 192–203.
- [47] E. Shi, T.-H. H. Chan, E. Rieffel, R. Chow, and D. Song, “Privacy-preserving aggregation of time-series data,” in *Network and Distributed System Security Symposium (NDSS)*, 2011.
- [48] R. Shokri and V. Shmatikov, “Privacy-preserving deep learning,” in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 2015, pp. 1310–1321.
- [49] R. Spring, A. Kyrillidis, V. Mohan, and A. Shrivastava, “Compressing gradient optimizers via count-sketches,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 5946–5955.
- [50] C. Sun, Y. Fu, J. Zhou, and H. Gao, “Personalized privacy-preserving frequent itemset mining using randomized response,” *The Scientific World Journal*, vol. 2014, 2014.
- [51] D. P. Team, “Learning with privacy at scale.” [Online]. Available: <https://docs-assets.developer.apple.com/ml-research/papers/learning-with-privacy-at-scale.pdf>
- [52] N. Wang, X. Xiao, Y. Yang, J. Zhao, S. C. Hui, H. Shin, J. Shin, and G. Yu, “Collecting and analyzing multidimensional data with local differential privacy,” in *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, April 2019, pp. 638–649.
- [53] S. Wang, L. Huang, P. Wang, Y. Nie, H. Xu, W. Yang, X. Li, and C. Qiao, “Mutual information optimally local private discrete distribution estimation,” *CoRR*, vol. abs/1607.08025, 2016.
- [54] T. Wang, J. Blocki, N. Li, and S. Jha, “Locally differentially private protocols for frequency estimation,” pp. 729–745, 2017.
- [55] T. Wang, N. Li, and S. Jha, “Locally differentially private frequent itemset mining,” in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 127–143.
- [56] S. L. Warner, “Randomized response: A survey technique for eliminating evasive answer bias,” *Journal of the American Statistical Association*, 1965.
- [57] K. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg, “Feature hashing for large scale multitask learning,” in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 1113–1120.
- [58] M. Ye and A. Barg, “Optimal schemes for discrete distribution estimation under locally differential privacy,” *IEEE Transactions on Information Theory*, vol. 64, no. 8, pp. 5662–5676, 2018.
- [59] Q. Ye, H. Hu, X. Meng, and H. Zheng, “Privkv: Key-value data collection with local differential privacy,” in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 317–331.
- [60] Q. Ye, H. Hu, X. Meng, H. Zheng, K. Huang, C. Fang, and J. Shi, “Privkvm*: Revisiting key-value statistics estimation with local differential privacy,” *IEEE Transactions on Dependable and Secure Computing*, 2021.
- [61] S. Yuan, M. Shen, I. Mironov, and A. C. A. Nascimento, “Practical, label private deep learning training based on secure multiparty computation and differential privacy,” *IACR Cryptol. ePrint Arch.*, p. 835, 2021.
- [62] Y. Zhao, J. Zhao, M. Yang, T. Wang, N. Wang, L. Lyu, D. Niyato, and K.-Y. Lam, “Local differential privacy-based federated learning for internet of things,” *IEEE Internet of Things Journal*, vol. 8, no. 11, pp. 8836–8853, 2020.

- [63] Y. Zhao, Y. Zhang, Y. Li, Y. Zhou, C. Chen, T. Yang, and B. Cui, “Minmax sampling: A near-optimal global summary for aggregation in the wide area,” in *Proceedings of the 2022 International Conference on Management of Data*, 2022.

APPENDIX A ADDITIONAL PRELIMINARIES

Theorem 17 (Sequential Composition Theorem [26]). *Assume the distribution of X and X' are (ϵ_1, δ_1) -close. If for any of $x \in \text{Domain}(X)$, the posterior distribution of random variable Y conditioned on $X = x$ and random variable Y' conditioned on $X' = x$ are (ϵ_2, δ_2) -close, then the distribution of (X, Y) and (X', Y') are $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -close.*

Theorem 18 (Post Processing Theorem [26]). *Assume the distribution of X and X' are (ϵ, δ) -close. Then for any (randomized) function f , the distributions of $f(X)$ and $f(X')$ are (ϵ, δ) -close.*

APPENDIX B RELATED WORK

Frequency estimation under LDP. Privacy-preserving frequency estimation is a fundamental primitive in federated analytics. Earlier works in this space focused on the case when the universe size d is small, and these works often suffer from per-client communication cost proportional to d . For example, RAPPOR and its variants [27], [53], [58] encode each client’s item with one-hot encoding and performs coordinate-wise randomized response (RR) [56], which suffers from at least d communication cost. Various subsequent works [2], [3], [7], [8], [12], [18], [28], [53], [54] focused on how to compress the communication especially when the universe size d is large, but each client has only one non-zero coordinate (i.e., the 1-sparse case). Some of these algorithms [2], [7], [8], [12], [18], [54] achieved optimal estimation error and using only logarithmic bandwidth.

When each user can have up to k items, one approach is to ask users to sample one item to report (e.g., [46], [55]) using the 1-sparse protocol (reviewed above) as a black-box. This approach introduces an error that is a \sqrt{k} factor away from optimal for user-level LDP, and a k factor away from optimal for event-level LDP.

Vector mean estimation under LDP. For vector mean estimation under LDP, a few earlier works [23], [24]. showed how to achieve optimal error for the dense case when $d \approx k$, absent communication constraints. Bhowmick et al. [9] showed how to achieve asymptotically optimal accuracy when $\epsilon > 1$, but they require $\Omega(d)$ communication. Following works, such as Harmony [43], Wang et al. [52], Li et al. [37] and Zhao et al. [62] improve the utility compared to [23]. However, all of the above works focused on the dense case and did not consider sparse vectors. Chen et al. [18] achieved optimal error and succinct communication for the 1-sparse case. The PrivKVM work [59] proposed an *interactive* protocol for vector mean estimation but it suffers from at least \sqrt{d} error; the approach was later improved [60] but the protocol is still interactive.

Practical challenges with LDP. LDP mechanisms have been deployed in practice [27], but two concerns have been raised by practitioners. One is the privacy leakage of LDP in the repeated setting, where repeated aggregation processes are performed for the same client set and the privacy leakage accumulates over time. Memoization, used by Rappor [27], is one of the standard techniques to prevent this circumstance by reusing the random seed for the local randomizers, making the privacy leakage as a one-time event. Also see [14], [31] for other works on the release of continuous DP statistics.

Another potential issue is data poisoning attacks [13], where malicious clients contribute intentionally corrupt data to disrupt the aggregation results. Restricting individual clients’ contributions and clipping are the most common techniques to prevent data poisoning. Moreover, assuming a common reference string (of length equal to the security parameter), generic zero-knowledge proofs can also be used to prove that their input vector is within some anticipated range and the clients computed the local randomizer correctly. This limits the influence of each client. See [11], [21] for related discussions.

Sparse aggregation and data release in the central setting. Previous works [4], [20], [36], [63] discussed the differentially private release of a k -sparse vector $v \in [0, u]^d$ (u can be large) in the central setting. The definition of neighboring databases in this work is also defined by L_1 distance; a neighboring input pair $v \sim v'$ iff $\|v - v'\|_1 \leq 1$. For example, recent work in this area presents the ALP mechanism [4] and shows how to privately encode a k -sparse vector with $O(k \log(d + u))$ bits with L_∞ decoding error of $O(\frac{\log d}{\epsilon})$. However, these works are either biased [4], [36] or require the non-zero coordinates of the vector to be significant, i.e., $\Omega(\log d/\epsilon)$ [20]. As such, they are not directly applicable to our aggregation setting.

Other definitions of differential privacy. Rényi differential privacy (RDP) [40] is a generalization of (ϵ, δ) -differential privacy based on Rényi divergence, such that it allows smoother composition with multiple private mechanisms [29]. We leave designing RDP vector aggregation mechanisms as an open question.

An orthogonal line of work considered computational differential privacy (CDP) [41] in distributed analytics [10], [15], [47], [61]. Some of these works showed how to compute distributed summation with error comparable to central DP, relying on cryptographic assumptions. Recently, Bagdasaryan *et al* [5] considered frequency estimation under CDP assuming 1-sparsity, with the extra assumption that the frequency vector must be sparse too. For the more general setting of k -sparsity that we consider, it is not known how CDP can further improve the accuracy in comparison with LDP, while still preserving succinct communication. We leave this as an open question.

Count sketching. Our methodology is also similar to the well-known count sketch algorithm [16], a dimensionality reduction technique widely used in non-private settings. Count sketch has had numerous applications such as feature hashing [57] and other machine learning and statistical applications [49].

However, most of these works are not designed with privacy in mind. Our techniques for making the algorithm private also differ from techniques used in the non-private literature. For example, feature hashing relies on the insight the L_2 distance is preserved by the hashing with sufficiently many bins, but we instead rely on the insight that the L_1 distance is compressed by the hashing.

APPENDIX C
ADDITIONAL DETAILS OF OUR UPPER BOUND
CONSTRUCTION

A. Discretization of Real Values for Communication

The clients need to send the reports tuple $(h_i, s_i, \tilde{B}_{i,1}, \dots, tB_{i,b})$ to the server. For the h_i and s_i , the client can send the random seed for the PRF to the server and the communication cost is $O(\xi)$. Here, ξ is the security parameter and with only $\text{negl}(\xi)$ probability, the randomness will be broken. The bucket values are unbounded real values. We actually know the “raw bucket value” $B_{i,1}, \dots, B_{i,b}$ are trivially bounded by $[-k, k]$. The unbounded part comes from the Laplacian noise and it has good concentration property. We can clip the value again with range $[-U, U]$, where $U = k + (\frac{\lambda}{\epsilon}) \log \frac{10nb}{\beta}$. We know if a random variable $r \sim \text{Lap}(\frac{\lambda}{\epsilon})$, then $\Pr[|r| \geq t] \leq e^{-t\epsilon/\lambda}$. Using union bound, we know that the magnitudes of all nb Laplacian random variables are smaller than $(\frac{\lambda}{\epsilon}) \log \frac{10nb}{\beta}$ with prob. $1 - \beta/10$. Then, we know for all clients, the report values are bounded by $[-U, U]$ with prob. $1 - \beta/10$. Then, we discretize the value using the unbiased discretizer

$$\text{DSC}(x) = \begin{cases} \lfloor x \rfloor, & \text{w.p. } 1 - (x - \lfloor x \rfloor) \\ \lfloor x \rfloor + 1, & \text{w.p. } x - \lfloor x \rfloor \end{cases}$$

We need to prove the error introduced by discretization is small. We denote the discretized version of \tilde{B}_i^c as \mathbf{B}_i^c . Trivially, for $c \in [n]$, $|\tilde{B}_{h^c(x)} s^c(x) - \mathbf{B}_i^c s^c(x)| \leq 1$, and also $\mathbf{E}[\tilde{B}_{h^c(x)}] = \mathbf{E}[\mathbf{B}_i^c]$. Using Hoeffding’s inequality, we can prove that $|\frac{1}{n} \sum_{c \in [n]} \tilde{B}_{h^c(x)} s^c(x) - \frac{1}{n} \sum_{c \in [n]} \mathbf{B}_i^c s^c(x)| \leq O\left(\sqrt{\frac{\log(d/\beta)}{n}}\right)$ with prob. at least $1 - \beta/10d$. Combining the above argument, we conclude that error introduced by the communication process is asymptotically equal or less than the error introduced by other process.

B. Additional Details for the Privacy Proof

Claim 19 (Restatement of Claim13). *Given any two neighboring vectors v, v' . Taking the randomness of h and s , if $\Pr[\sum_{j \in [b]} |B_j - B'_j| \geq \lambda] \leq \delta$, then the two invocation of the local randomizers’ output distributions are (ϵ, δ) -close.*

Proof. The clipping process will only make the difference smaller. Conditioned on the case that $\sum_{j \in [b]} |B_j - B'_j| \leq \lambda$. Then, we can bound the ratio between the probability density

function of the r.v. $B_j + r_j$ and $B'_j + r_j$: for any $x_1, \dots, x_b \in \mathbb{R}$,

$$\begin{aligned} \frac{p_v(\bigwedge_{j \in [b]} B_j + r_j = x | h, s)}{p_{v'}(\bigwedge_{j \in [b]} B'_j + r_j = x | h, s)} &= \frac{\prod_{j \in [b]} \exp(\frac{\epsilon}{\lambda} \|x - B_j\|)}{\prod_{j \in [b]} \exp(\frac{\epsilon}{\lambda} \|x - B'_j\|)} \\ &= \exp\left(\frac{\epsilon}{\lambda} \sum_{j \in [b]} (|x - B_j| - |x - B'_j|)\right) \\ &\leq \exp\left(\frac{\epsilon}{\lambda} \sum_{j \in [b]} |B_j - B'_j|\right) \leq \exp(\epsilon). \end{aligned}$$

Therefore, the distribution of \tilde{B}_i and \tilde{B}'_i is $(\epsilon, 0)$ close. By post-processing theorem, the joint distribution of the final report $\tilde{B}_1, \dots, \tilde{B}_b$ and $\tilde{B}'_1, \dots, \tilde{B}'_b$ is still $(\epsilon, 0)$ -close. Considering the failure probability δ such that some randomly sampled h and s cause $\sum_{j \in [b]} |B_j - B'_j| > \delta$, the distributions for the whole outputs $(h, s, \tilde{B}_1, \dots, \tilde{B}_b)$ and $(h, s, \tilde{B}'_1, \dots, \tilde{B}'_b)$ are (ϵ, δ) -close. \square

C. Full Proof of the Utility Theorem

Below we give the full proof of the utility statement in Theorem 12.

Proof. Our proof bound the absolute error incurred step by step. Fix an index $x \in [d]$. The server computes the estimation \hat{v}_x as $\hat{v}_x = \frac{1}{n} \sum_{i \in [n]} s_i(x) \tilde{B}_{i, h_i(x)}$. We bound the error by the three steps: 1) binning; 2) clipping; 3) adding Laplacian noise.

$$\begin{aligned} &|\bar{v}_x - \hat{v}_x| \\ &= \left| \frac{1}{n} \sum_{i \in [n]} v_{i,x} - \frac{1}{n} \sum_{i \in [n]} s_i(x) \tilde{B}_{i, h_i(x)} \right| \\ &\leq \left| \frac{1}{n} \sum_{i \in [n]} v_{i,x} - \frac{1}{n} \sum_{i \in [n]} B_{i, h_i(x)} s_i(x) \right| \\ &\quad + \left| \frac{1}{n} \sum_{i \in [n]} B_{i, h_i(x)} s_i(x) - \frac{1}{n} \sum_{i \in [n]} \bar{B}_{i, h_i(x)} s_i(x) \right| \\ &\quad + \left| \frac{1}{n} \sum_{i \in [n]} \bar{B}_{i, h_i(x)} s_i(x) - \frac{1}{n} \sum_{i \in [n]} \tilde{B}_{i, h_i(x)} s_i(x) \right| \end{aligned}$$

We now look at the binning error term $|\frac{1}{n} \sum_{i \in [n]} v_{i,x} - \frac{1}{n} \sum_{i \in [n]} B_{i, h_i(x)} s_i(x)|$. We have $B_{i, h_i(x)} = \sum_{l \in [d]} \mathbb{I}[h_i(l) = h_i(x)] v_{i,l} s_i(l)$. Define random variables $Y_{i,l} = \mathbb{I}[h_i(l) = h_i(x)] v_{i,l} s_i(l) s_i(x)$ for $i \in [n], l \in [d]$. Then the error term can be rewritten as $|\frac{1}{n} \sum_{i \in [n]} \sum_{l \in [d], l \neq x} Y_{i,l}|$. Since $s_i(\cdot)$ is a random ± 1 function, $s_i(l) s_i(x)$ can be seen as an independent uniform ± 1 random variable. Also, the hash function h_i is a random oracle, so for $l \neq x$, $\Pr[\mathbb{I}[h_i(l) = h_i(x)] = 1] = \frac{1}{b}$.

So we know $\mathbf{E}[Y_{i,l}] = 0, \mathbf{E}[(Y_{i,l})^2] = \frac{1}{b}v_{i,l}^2 \leq \frac{1}{b}$. Thus, we can use the Bernstein's Inequality:

$$\begin{aligned} & \Pr \left[\left| \frac{1}{n} \sum_{i \in [n]} \sum_{l \in [d], l \neq x} Y_{i,l} \right| \geq t \right] \\ & \leq 2 \exp \left(- \frac{(nt)^2}{\sum_{i \in [n]} \sum_{l \in [d], l \neq x} \mathbf{E}[(Y_{i,l})^2] + \frac{1}{3}nt} \right) \\ & \leq 2 \exp \left(- \frac{(nt)^2}{\frac{nk}{b} + \frac{1}{3}nt} \right) = 2 \exp \left(- \frac{nt^2}{\frac{k}{b} + \frac{1}{3}t} \right). \end{aligned}$$

Using the fact that $k/b \geq \log(5d/\beta)/n$ and setting $t = \Theta(\sqrt{\frac{k}{b}} \sqrt{\frac{\log(d/\beta)}{n}})$ with a proper constant, we can prove that $\Pr[\frac{1}{n} \sum_{i \in [n]} \sum_{l \in [d], l \neq x} Y_{i,l} \geq t] \leq \beta/10d$. That means, $|\frac{1}{n} \sum_{i \in [n]} v_{i,x} - \frac{1}{n} \sum_{i \in [n]} B_{i,h_i(x)} s_i(x)| = O(\sqrt{\frac{k}{b}} \sqrt{\frac{\log(d/\beta)}{n}})$ with probability $1 - \beta/10d$.

Next, we look at the clipping error term $|\frac{1}{n} \sum_{i \in [n]} B_{i,h_i(x)} s_i(x) - \frac{1}{n} \sum_{i \in [n]} \bar{B}_{i,h_i(x)} s_i(x)|$. We actually try to prove the clipping range is large enough, so that, with probability $1 - \beta/2$, for all $i \in [n], j \in [b]$, $B_{i,j} = \bar{B}_{i,j}$, i.e., $|B_{i,j}| \leq \eta$. Then, the error term becomes zero naturally. Fix any i, j . $B_{i,j} = \sum_{l \in [d]} \mathbb{I}[h_i(l) = j] s_i(l) v_{i,l}$. We define a random variable $X_l = \mathbb{I}[h_i(l) = j] s_i(l) v_{i,l}$. Then we know $B_{i,j} = \sum_{l \in [d]} X_l$. Its moment generating function is

$$\begin{aligned} & \mathbf{E}[\exp(sB_{i,j})] = \prod_{l \in [d]} \exp(sX_l) \\ & = \prod_{l \in [d]} \left(1 - \frac{1}{b} + \frac{1}{2b} (e^{-sv_{i,l}} + e^{sv_{i,l}}) \right) \\ & \leq \prod_{l \in [d]} \frac{1}{2} (e^{-sv_{i,l}} + e^{sv_{i,l}}) \leq \prod_{l \in [d]} \exp(s^2 v_{i,l}^2 / 2) \\ & = \exp \left(\frac{s^2}{2} \sum_{l \in [d]} v_{i,l}^2 \right) \leq \exp \left(\frac{ks^2}{2} \right). \end{aligned}$$

Hence, $B_{i,j}$ is a sub-Gaussian random variable with variance k . When $\eta \geq \sqrt{2k \log(4nb/\beta)}$, $\Pr[|B_{i,j}| \geq \eta] \leq 2 \exp(-\frac{\eta^2}{2k}) = 2 \exp(-\log(4nb/\beta)) = \beta/2nb$. Using union bound, we prove that, with prob. $1 - \beta/2$, for all $i \in [n], j \in [b]$, $|B_{i,j}| \leq \eta$, i.e., $\bar{B}_{i,j} = B_{i,j}$.

Now, we look at the error term, $|\frac{1}{n} \sum_{i \in [n]} \bar{B}_{i,h_i(x)} s_i(x) - \frac{1}{n} \sum_{i \in [n]} \tilde{B}_{i,h_i(x)} s_i(x)|$. We know that for all $j \in [b], i \in [n]$, $\bar{B}_{i,j} = \tilde{B}_{i,j} + r_{i,j}$, where $r_{i,j} \sim \text{Lap}(\frac{\lambda}{\epsilon})$. Also, using the fact that Laplacian distribution is symmetrical over positive value and negative value and $s_i(x)$ is a uniform random ± 1 variable, the error term can be rewritten as $|\frac{1}{n} \sum_{i \in [n]} r_{i,h_i(x)}|$. We know that the Laplacian noise $\text{Lap}(\frac{\lambda}{\epsilon})$'s distribution is a sub-exponential distribution $\text{subE}(\frac{\lambda}{\epsilon})$. Using the concentration bound for sub-exponential random variable, we can prove that with prob. $1 - \beta/10d$, $|\frac{1}{n} \sum_{i \in [n]} \bar{B}_{i,h_i(x)} s_i(x) - \frac{1}{n} \sum_{i \in [n]} \tilde{B}_{i,h_i(x)} s_i(x)| \leq O\left(\frac{\lambda}{\epsilon} \sqrt{\frac{\log(d/\beta)}{n}}\right)$.

Taking the union bound over all $x \in [d]$, we prove that with probability at least $1 - 5\beta/4$, $\max_{x \in [d]} |\frac{1}{n} \sum_{i \in [n]} v_{i,x} - \frac{1}{n} \sum_{i \in [n]} s_i(x) \tilde{B}_{i,h_i(x)}| \leq O\left(\left(\sqrt{\frac{k}{b}} + \frac{\lambda}{\epsilon}\right) \sqrt{\frac{\log(d/\beta)}{n}}\right)$. \square

APPENDIX D DETAILED LOWER BOUND PROOF

In this section, we give the detailed proof of Theorem 15.

Estimation of distribution mean. Because the empirical average \bar{v} is concentrated around the distribution mean $\mathbf{E}[\mathcal{P}]$ (Lemma 20), it suffices to consider the expected error of estimating $\mathbf{E}[\mathcal{P}]$ by the following quantity:

$$\mathcal{E}(\mathcal{A}; \mathcal{P}) \triangleq \mathbf{E}[\|\mathcal{A}(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n) - \mathbf{E}[\mathcal{P}]\|_\infty], \quad (3)$$

where the randomness comes from sampling v_i from \mathcal{P} , the randomized algorithms \mathcal{Q}_i 's from all clients and the estimator \mathcal{A} . The following result (which is also used in [8]) implies that it suffices to prove the same asymptotic lower bound for $\mathcal{E}(\mathcal{A}; \mathcal{P})$ to achieve Theorem 15.

Lemma 20 (Empirical Average vs Distribution Mean). *Let \bar{V} be the empirical average of n i.i.d. samples from \mathcal{P} . Then, $\mathbf{E}[\|\bar{V} - \mathbf{E}[\mathcal{P}]\|_\infty] \leq O\left(\sqrt{\frac{\log d}{n}}\right) = o\left(\frac{1}{\epsilon} \sqrt{\frac{k \log(d/k)}{n}}\right)$.*

Analyzing expected error via an encoding-decoding process. Given randomized algorithms \mathcal{Q}_i 's and aggregator function \mathcal{A} , the lower bound framework in [8] consider the following encoding-decoding process. Denote $\eta := \min\{\Theta(\frac{1}{\epsilon} \sqrt{\frac{\log |\mathcal{S}|}{n}}), 1\}$.

- 1) Sample V uniformly at random from \mathcal{S} .
- 2) Each client i receives the same V from the previous step, and performs the following actions independently.
 - Sample v_i from $\mathcal{P}_V^{(\eta)}$, where for $v \in \mathcal{S}$, the distribution $\mathcal{P}_v^{(\eta)}$ is defined as in (2).
 - Apply local LDP mechanism $\mathcal{Q}_i(\cdot)$ to obtain $\mathbf{z}_i := \mathcal{Q}_i(v_i)$.
- 3) Using the aggregator function $\mathcal{A}(\cdot)$, compute $Y = \frac{1}{\eta} (\mathcal{A}(\mathbf{z}_1, \dots, \mathbf{z}_n) - (1 - \eta) \cdot \frac{1}{n} \sum_{v \in \mathcal{S}} v)$. Round Y to $\hat{V} \in \{0, 1\}^d$, i.e., for each $j \in [d]$, $\hat{V}_j := 1$ if $Y_j \geq \frac{1}{2}$, and 0 otherwise.
- 4) Define the event error as $V \neq \hat{V}$.

The crux of the proof depends on the following bounds on $\Pr[\text{error}]$:

- A lower bound by Fano's Inequality:

$$\Pr[\text{error}] \geq 1 - \frac{I(V; \mathbf{z}_1, \dots, \mathbf{z}_n) + 1}{\log |\mathcal{S}|}. \quad (4)$$

Since conditioning on V , the \mathbf{z}_i 's are independent, we have: $I(V; \mathbf{z}_1, \dots, \mathbf{z}_n) = \sum_{i \in [n]} I(V; \mathbf{z}_i)$.

An upper bound on $I(V; \mathbf{z}_i) = I(V; \mathcal{Q}_i(v_i))$ using the differential privacy of \mathcal{Q}_i will be given in Lemmas 22 and 23.

- An upper bound on $\Pr[\text{error}]$ is given in Lemma 21.

Lemma 21 (Low Decoding Error). *Suppose for all $v \in \mathcal{S}$, $\mathcal{E}(\mathcal{A}; \mathcal{P}_v^{(\eta)}) \leq \frac{\eta}{10}$. Then, $\Pr[\text{error}] \leq \frac{1}{5}$.*

Proof. Observe that the event error implies that for at least one $j \in [d]$, the difference between the j -th coordinates of Y and V is at least $\frac{1}{2}$, i.e., $\|Y - V\|_\infty \geq \frac{1}{2}$.

Hence, by Markov's inequality, the probability of this event is at most $2 \cdot \mathbf{E}[\|Y - V\|_\infty]$. Finally, as shown in [8], observe that: $\mathbf{E}[\|Y - V\|_\infty] = \frac{1}{\eta} \cdot \mathbf{E}_V[\mathcal{E}(\mathcal{A}; \mathcal{P}_V^{(\eta)})] \leq \frac{1}{10}$, which gives the result. \square

Bounding mutual information via differential privacy. The following lemmas from [8] give upper bounds on the mutual information between the input and the output of differentially private algorithms.

Lemma 22. *Suppose $0 < \epsilon = O(1)$ and $0 < \delta < \epsilon$. Let V be a random variable that is uniformly distributed on a discrete set \mathcal{S} . Suppose the output of the (randomized) algorithm $\mathcal{Q} : \mathcal{S} \rightarrow \mathcal{Z}$ is (ϵ, δ) -differentially private where any two inputs in \mathcal{S} are considered as neighboring. Then, the mutual information between the input variable and the output report is bounded:*

$$I(V; \mathcal{Q}(V)) = O(\epsilon^2 + \frac{\delta}{\epsilon} \log |\mathcal{S}| + \frac{\delta}{\epsilon} \log(\epsilon/\delta))$$

Lemma 23. *Suppose $0 < \epsilon = O(1)$ and $0 < \delta < 1$ and $\mathcal{Q} : \mathcal{S} \rightarrow \mathcal{Z}$ is (ϵ, δ) -differentially private. Define the algorithm $\mathcal{Q}^{(\eta)} : \mathcal{S} \rightarrow \mathcal{Z}$ as follows: on input $v \in \mathcal{S}$, sample V from $\mathcal{P}_v^{(\eta)}$ (which is defined in the encoding-decoding procedure) and return $\mathcal{Q}(V)$. Then, the output of $\mathcal{Q}^{(\eta)}$ is $(O(\eta\epsilon), O(\eta\delta))$ -differentially private.*

Finalizing the proof of Theorem 15. For the sake of contradiction, we assume that for any distribution \mathcal{P} on \mathcal{S} , $\mathcal{E}(\mathcal{A}; \mathcal{P}) \leq \frac{\eta}{10}$. Then, Lemma 21 implies that decoding error happens with $\Pr[\text{error}] \leq \frac{1}{5}$.

In view of Fano's Inequality (4), a contradiction can be achieved if $\frac{I(V; \mathbf{z}_1, \dots, \mathbf{z}_n) + 1}{\log |\mathcal{S}|} \leq \frac{1}{2}$.

By Lemmas 22 and 23, for each i , $I(V; \mathbf{z}_i) \leq O(\eta^2 \epsilon^2 + \frac{\delta}{\epsilon} \log |\mathcal{S}| + \frac{\delta}{\epsilon} \log(\epsilon/\delta))$.

By choosing sufficiently small $\eta = \min\{\Theta(\frac{1}{\epsilon} \sqrt{\frac{\log |\mathcal{S}|}{n}}, 1)\}$ and $\delta = o(\frac{\epsilon}{n \log n})$, it follows that:

$\frac{I(V; \mathbf{z}_1, \dots, \mathbf{z}_n) + 1}{\log |\mathcal{S}|} = \frac{\sum_{i \in [n]} I(V; \mathbf{z}_i) + 1}{\log |\mathcal{S}|} < \frac{1}{2}$, where the first equality holds because conditioning on V , the \mathbf{z}_i 's are independent.

Hence, we have obtained the desired contradiction that completes the proof of Theorem 15.