Heavy Hitters and the Structure of Local Privacy

Mark Bun
Princeton University
mbun@cs.princeton.edu

Jelani Nelson*
Harvard University
minilek@seas.harvard.edu

Uri Stemmer†
Weizmann Institute
u@uri.co.il

ABSTRACT

We present a new locally differentially private algorithm for the heavy hitters problem which achieves optimal worst-case error as a function of all standardly considered parameters. Prior work obtained error rates which depend optimally on the number of users, the size of the domain, and the privacy parameter, but depend sub-optimally on the failure probability.

We strengthen existing lower bounds on the error to incorporate the failure probability, and show that our new upper bound is tight with respect to this parameter as well. Our lower bound is based on a new understanding of the structure of locally private protocols. We further develop these ideas to obtain the following general results beyond heavy hitters.

- *Advanced Group Privacy:* In the local model, group privacy for k users degrades proportionally to $\approx \sqrt{k}$, instead of linearly in k as in the central model. Stronger group privacy yields improved max-information guarantees, as well as stronger lower bounds (via “packing arguments”), over the central model.
- Building on a transformation of Bassily and Smith (STOC 2015), we give a generic transformation from any non interactive approximate-private local protocol into a pure-private local protocol. Again in contrast with the central model, this shows that we cannot obtain more accurate algorithms by moving from pure to approximate local privacy.

KEYWORDS

Differential privacy, heavy hitters, local model

ACM Reference Format:

Mark Bun, Jelani Nelson, and Uri Stemmer. 2018. Heavy Hitters and the Structure of Local Privacy. In *PODS'18: 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, June 10–15, 2018, Houston, TX, USA*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3196959.3196981>

*Supported by NSF grant IIS-1447471 and CAREER award CCF-1350670, ONR Young Investigator award N00014-15-1-2388 and DORECG award N00014-17-1-2127, an Alfred P. Sloan Research Fellowship, and a Google Faculty Research Award.

†Work done while the third author was a postdoctoral fellow at the Center for Research on Computation and Society, Harvard University, supported by NSF grant 1565387.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

PODS'18, June 10–15, 2018, Houston, TX, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-4706-8/18/06...\$15.00

<https://doi.org/10.1145/3196959.3196981>

1 INTRODUCTION

In the *heavy-hitters problem*, each of n users holds as input an item from some domain X . Our goal is to identify all “heavy-hitters”, which are the domain elements $x \in X$ for which many of the users have x as their input. In this work we study the heavy-hitters problem in the *local model* of differential privacy (LDP), where the users randomize their data locally, and send differentially private reports to an untrusted server that aggregates them. In our case, this means that every user only sends the server a single randomized message that leaks very little information about the input of that user.

The heavy hitters problem is perhaps the most well-studied problem in local differential privacy [2, 3, 14, 18, 21]. In addition to the intrinsic interest of this problem, LDP algorithms for heavy-hitters provide important subroutines for solving many other problems, such as median estimation, convex optimization, and clustering [20, 24]. In practice, heavy-hitters algorithms under LDP have already been implemented to provide organizations with valuable information about their user bases while providing strong privacy guarantees to their users. Two prominent examples are by Google in the Chrome browser [11] and Apple in iOS-10 [25], making local heavy-hitter algorithms the most widespread industrial application of differential privacy to date.

Before describing our new results, we define our setting more precisely. Consider a set of n users, where each user i holds as input an item $x_i \in X$. We denote $S = (x_1, \dots, x_n)$, and refer to S as a “distributed database” (as it is not stored in one location, and every x_i is only held locally by user i). We say that a domain element $x \in X$ is Δ -heavy if its multiplicity in S is at least Δ , i.e., if there are at least Δ users whose input is x . Our goal is to identify all Δ -heavy elements for Δ as small as possible. Specifically, we receive a sequence of messages—one message y_i from every user i —at the end of which our goal is to report a list $L \subset X$ such that

- (1) Every Δ -heavy element x is in L , and
- (2) $|L| = O(n/\Delta)$, i.e., the list size is proportional to the maximum possible number of Δ -heavy elements.

We also want, for each $x \in L$, to output a value which is within Δ of the multiplicity of x in S . The parameter Δ is referred to as *the error* of the protocol, as elements with multiplicities less than Δ are “missed” by the protocol. We refer to the probability that the protocol does not successfully achieve the above objectives as its *failure probability*, denoted by β . The privacy requirement is that for every user i , the distribution of the message y_i should not be significantly affected by the input of the user. Formally,

DEFINITION 1.1 ([8]). An algorithm $\mathcal{A} : X \rightarrow Y$ satisfies ϵ -differential privacy if for every two inputs x, x' and for every output $y \in Y$ we have

$$\Pr[\mathcal{A}(x) = y] \leq e^\epsilon \cdot \Pr[\mathcal{A}(x') = y].$$

Solving the heavy-hitters problem under LDP becomes harder for smaller choices of Δ ,¹ and the work on LDP heavy-hitters have focused on constructing protocols that can achieve as small a value of Δ as possible.

In this work we combine ideas from the recent locally-private heavy-hitters algorithm of [2] with the recent non-private algorithm of [17], and present a new efficient LDP heavy-hitters algorithm achieving optimal worst-case error Δ as a function of the number of users n , the size of the domain $|X|$, the privacy parameter ϵ , and the failure probability β . Prior state-of-the-art results [2, 3] either obtained error with sub-optimal dependence on the failure probability β , or had runtime at least linear in $|X|$ (which is not realistic in real applications, where X may be the space of all reasonable-length URL domains). Table 1 summarizes our result in comparison with previous algorithms for the problem.

We remark that while our focus is on achieving the optimal (i.e., smallest possible) error Δ , our algorithm also allows the server to obtain larger error $\Delta' \geq \Delta$ in exchange for lower space requirements. Specifically, memory usage and the size of the output list may both be improved to $\tilde{O}(n/\Delta')$.

1.1 Lower Bound on the Error and New Understandings of the Local Model

We strengthen a lower bound on the error by [5, 14, 26] and [3] to incorporate the failure probability, and show that our new upper bound is tight. Our lower bound is based (conceptually) on the following new understandings of the local privacy model:

1.1.1 Advanced Group Privacy. An important property of differential privacy is *group privacy*. Recall that differential privacy ensures that no single individual has a significant effect on the outcome (distribution) of the computation. A similar guarantee also holds for every small *group* of individuals, but the privacy guarantee degrades (gracefully) with the size of the group. Specifically, if an algorithm satisfies an individual-level guarantee of ϵ -differential privacy, then it satisfies $k\epsilon$ -differential privacy for every group of k individuals. We observe that in the local model, group privacy degrades the privacy parameter only by a factor of $\approx \sqrt{k}$, unlike in the centralized model. This observation allows us to show improved bounds on the *max-information*—a measure of the dependence of an algorithm's output on a randomly chosen input—of locally-private protocols [7]. Strong group privacy is, however, a mixed blessing as it also leads to stronger lower bounds (of the type known as “packing arguments”) for pure-private local algorithms.

1.1.2 Pure vs. Approximate LDP. Another important property of differential privacy is that it also degrades gracefully under *composition*. This allows us to design algorithms that access their input database using several (say k) differentially private mechanisms as subroutines, and argue about the overall privacy guarantees. In the case of pure-privacy, the privacy parameter deteriorates by a factor of at most k . However, in the case of approximate-privacy, k -fold

composition degrades the privacy parameter only by a factor of $\approx \sqrt{k}$ (see [10]).

We show that in the non-interactive local model we can enjoy the best of both worlds. First, we can show that in some cases, (an approximate version of) the composition of locally private algorithms can satisfy *pure*-privacy while enjoying the same strong composition guarantees of approximate-privacy. Motivated by this observation, we then proceed to show that in the local model approximate-privacy is actually never more useful than pure-privacy (at least for non-interactive protocols). Specifically, building on a transformation of Bassily and Smith [3], we give a generic transformation from any non-interactive approximate-private local protocol into a pure-private local protocol with the same utility guarantees. This is the first formal proof that approximate local-privacy cannot enable more accurate analyses than pure local-privacy.

In Section 7 we describe how our new lower bound for the heavy hitters under *pure*-LDP follows from advanced group privacy, and how the lower bound extends to *approximate*-LDP via our generic transformation. In addition, we provide a direct analysis for our lower bound that results from a direct extension of the argument of [5, 14, 26] (using ideas from [3]).

2 PRELIMINARIES FROM DIFFERENTIAL PRIVACY

Notations. Databases are (ordered) collections of elements from some data universe X . Two databases $S, S' \in X^n$ are called *neighboring* if they differ in at most one entry. Throughout this paper, we use the notation $\tilde{O}$ to hide logarithmic factors in n and $|X|$ (the number of users and the size of the domain).

Consider a database where each entry contains information pertaining to one individual. An algorithm operating on databases is said to preserve differential privacy if a change of a single record of the database does not significantly change the output distribution of the algorithm. Intuitively, this means that individual information is protected: whatever is learned about an individual could also be learned with her data arbitrarily modified (or without her data at all).

DEFINITION 2.1 ([8]). A randomized algorithm $M : X^n \rightarrow Y$ is (ϵ, δ) differentially private if for every two neighboring datasets $S, S' \in X^n$ and every $T \subseteq Y$ we have $\Pr[M(S) \in T] \leq e^\epsilon \cdot \Pr[M(S') \in T] + \delta$, where the probability is over the randomness of M .

2.1 Local Differential Privacy

In the local model, each individual holds her private information locally and only releases the outcomes of privacy-preserving computations on her own data. This is modeled by letting the algorithm access each entry x_i in the input database $S = (x_1, \dots, x_n) \in X^n$ separately, via differentially private *local randomizers*.

DEFINITION 2.2 ([8, 16]). A local randomizer $R : X \rightarrow W$ is a randomized algorithm that takes a database of size $n = 1$. Let $S = (x_1, \dots, x_n) \in X^n$ be a database. An LR oracle $LR_S(\cdot, \cdot)$ gets an index $i \in [n]$ and a local randomizer R , and outputs a random value $w \in W$ chosen according to the distribution $R(x_i)$.

¹For intuition, observe that we cannot hope to have $\Delta = 1$. Indeed, if only one user is holding some x^* as input (so x^* is 1-heavy), then the outcome distribution of the protocol remains approximately the same even when this user holds a different input (in which case x^* does not appear in the data at all, and it is no longer 1-heavy).

Performance metric	This work	Bassily et al. [2]	Bassily and Smith [3] ²
Server time	$\tilde{O}(n)$	$\tilde{O}(n)$	$\tilde{O}(n^{2.5})$
User time	$\tilde{O}(1)$	$\tilde{O}(1)$	$\tilde{O}(n^{1.5})$
Server memory	$\tilde{O}(\sqrt{n})$	$\tilde{O}(\sqrt{n})$	$\tilde{O}(n^2)$
User memory	$\tilde{O}(1)$	$\tilde{O}(1)$	$\tilde{O}(n^{1.5})$
Communication/user	$O(1)$	$O(1)$	$O(1)$
Public randomness/user	$\tilde{O}(1)$	$\tilde{O}(1)$	$\tilde{O}(n^{1.5})$
Worst-case error	$O\left(\frac{1}{\varepsilon} \cdot \sqrt{n \log\left(\frac{ X }{\beta}\right)}\right)$	$O\left(\frac{1}{\varepsilon} \cdot \sqrt{n \log\left(\frac{ X }{\beta}\right) \log\left(\frac{1}{\beta}\right)}\right)$	$O\left(\frac{\log^{1.5}\left(\frac{1}{\beta}\right)}{\varepsilon} \cdot \sqrt{n \log X }\right)$

Table 1: Comparison with previous works. The $\tilde{O}$ notation hides logarithmic factors in n and $|X|$.

DEFINITION 2.3 ([8, 16]). *An algorithm satisfies (ε, δ) -local differential privacy (LDP) if it accesses the database S only via the oracle LR_S with the following restriction: for all possible executions of the algorithm and for all $i \in [n]$, if $LR_S(i, R_1), \dots, LR_S(i, R_k)$ are the algorithm's invocations of LR_S on index i , then the algorithm $\mathcal{B}(x) = (R_1(x), R_2(x), \dots, R_k(x))$ is (ε, δ) -differentially private.*

In the above definition, k is the number of invocations of LR_S on index i made throughout the execution (potentially this number can be different for different users, and can even be chosen adaptively throughout the execution). In our heavy-hitters protocol we will have $k = 1$, meaning that every user only sends one response to the server. Local algorithms that prepare all their queries to LR_S before receiving any answers are called *non-interactive*; otherwise, they are *interactive*. The focus of this work is on non-interactive protocols.

3 A HEAVY-HITTERS PROTOCOL WITH OPTIMAL ERROR

Let $S = (x_1, \dots, x_n) \in X^n$ be a database, which may be distributed across n users (each holding one row). For $x \in X$, we will be interested in estimating the frequency (or multiplicity) of x in S , denoted as $f_S(x) = |\{i \in [n] : x_i = x\}|$. Specifically, we would like to solve the following problem while guaranteeing LDP:

DEFINITION 3.1. *Let $S \in X^n$ be a (distributed) database. In the heavy-hitters problem with error Δ and failure probability β , the goal is to output a list $\text{Est} \subseteq X \times \mathbb{R}$ of elements and estimates, of size $|\text{Est}| = O(n/\Delta)$, such that with probability $1 - \beta$*

- (1) *For every $(x, a) \in \text{Est}$ we have $|f_S(x) - a| \leq \Delta$; and,*
- (2) *$\forall x \in X$ s.t. $f_S(x) \geq \Delta$ we have that x appears in Est .*

As in previous works, our algorithm for the heavy-hitters problem is based on a reduction to the simpler task of constructing a *frequency oracle*:

DEFINITION 3.2. *Let $S \in X^n$ be a (distributed) database. A frequency oracle with error Δ and failure probability β is an algorithm*

that, with probability at least $1 - \beta$, outputs a data structure capable of approximating $f_S(x)$ for every $x \in X$ with error at most Δ .

Observe that constructing a frequency oracle is an easier task than solving the heavy-hitters problem, as every heavy-hitters algorithm is in particular a frequency oracle. Specifically, given a solution Est to the heavy-hitters problem, we can estimate the frequency of every $x \in X$ as $\hat{f}_S(x) = a$ if $(x, a) \in \text{Est}$, or $\hat{f}_S(x) = 0$ otherwise. Moreover, observe that ignoring runtime, the two problems are identical (since we can query the frequency oracle on every domain element to find the heavy-hitters).

The literature on heavy-hitters and frequency oracles under LDP has focused on the goal of minimizing the error Δ . However, as Bassily and Smith showed, under LDP, every algorithm for either task must have worst case error at least $\Omega\left(\frac{1}{\varepsilon} \sqrt{n \cdot \log |X|}\right)$.

This lower bound is known to be tight, as constructions of frequency oracles with matching error were presented in several works [2, 3, 14, 18]. These works also presented reductions that transform a frequency oracle into a heavy-hitters algorithm. However, their reductions resulted in a sub-optimal dependency of the error in the failure probability β (on which the lower bound above is not informative). In this work we give a new (privacy preserving) reduction from the heavy-hitters problem to the frequency oracle problem, achieving error which is optimal in all parameters. Our reduction is a private variant of the recent non-private algorithm of [17], and our final construction uses the frequency oracle of [2].

3.1 Existing Techniques

In this section we give an informal survey of the techniques of [2] and [17], and highlight some of the ideas behind the constructions. This intuitive overview is generally oversimplified, and hides many of the difficulties that arise in the actual analyses.

3.1.1 Reduction with Sub-Optimal Dependence on the Failure Probability. Assume that we have a frequency oracle protocol with worst-case error Δ , and let $S = (x_1, \dots, x_n) \in X^n$ be a (distributed) database, where user i holds the input x_i . We now want to use our frequency oracle to identify a small set of “potential heavy-hitters”. Specifically, we would like to identify all elements whose frequencies in S are at least 2Δ . We outline the approach of [2].

²The user's run-time and memory in [3] can be improved to $O(n)$ if one assumes random access to the public randomness, which we do not assume in this work. In fact, our protocol can be implemented without public randomness while attaining essentially the same performances.

Let $h : X \rightarrow [Y]$ be a (publicly known) random hash function, mapping domain elements into some range of size Y . We will now use h in order to identify the heavy-hitters. To that end, let $x^* \in X$ denote such a heavy-hitter, appearing at least 2Δ times in the database S , and denote $y^* = h(x^*)$. Assuming that Y is big enough, w.h.p. we will have that x^* is the only input element (from S) that is mapped (by h) to the hash value y^* . Assuming that this is indeed the case, we will now attempt to identify x^* one “coordinate” (as defined below) at a time.

Let us assume that elements from X are represented as M symbols from an alphabet $[W]$. For every coordinate $m \in [M]$, denote $S_m = (h(x_i), x_i[m])_{i \in [n]}$, where $x_i[m]$ is the m^{th} symbol of x_i . That is, S_m is a database over the domain $([Y] \times [W])$, where the row corresponding to user i is $(h(x_i), x_i[m])$. Observe that since h is public, every user can compute her own row locally. As x^* is a heavy-hitter, for every $m \in [M]$ we have that $(y^*, x^*[m])$ appears in S_m at least 2Δ times. On the other hand, as we assumed that x^* is the only input element that is mapped to y^* , we get that (y^*, w') does not appear in S_m at all for every $w' \neq x^*[m]$. Recall that our frequency oracle has error at most Δ , and hence, we can use it to accurately determine the symbols of x^* . Using this strategy we can identify a set of size Y containing all heavy-hitters: For every hash value $y \in [Y]$, construct a potential heavy-hitter $\hat{x}^{(y)}$ symbol-by-symbol, where the m^{th} symbol of $\hat{x}^{(y)}$ is the symbol w s.t. (y, w) is most frequent in S_m . By the above discussion, this strategy identifies all of the heavy-hitters.

Recall that we assumed here that Y (the size of the range of the hash function h) is big enough so that there are no collisions among input elements (from S). In fact, the analysis can withstand a small number of collisions (roughly $\sqrt{n}$, since we are aiming for error bigger than $\sqrt{n}$ anyway). Nevertheless, for the above strategy to succeed with high probability, say w.p. $1 - \beta$, we will need to set $Y \gtrsim \sqrt{n}/\beta$. As β can be exponentially small, this is unacceptable, and hence, Bassily et al. [2] applied the above strategy with a large failure probability β , and amplified the success probability using repetitions. Unfortunately, using repetitions still causes the error to increase by a factor of $\sqrt{\log(1/\beta)}$ (because applying multiple private computations to the same dataset degrades privacy). Specifically, Bassily et al. [2] obtained the following result.

THEOREM 3.3 ([2]). *Let $\beta, \epsilon \leq 1$, and let $S \in X^n$ be a (distributed) database. There exists an ϵ -LDP algorithm that returns a list Est of length $\tilde{O}(\sqrt{n})$ such that with probability $1 - \beta$*

- (1) $\forall (x, a) \in \text{Est}, |a - f_S(x)| \leq O\left(\frac{1}{\epsilon} \sqrt{n \log\left(\frac{\min\{n, |X|\}}{\beta}\right)}\right)$.
- (2) $\forall x \in X$ s.t. $f_S(x) \geq O\left(\frac{1}{\epsilon} \sqrt{n \log\left(\frac{|X|}{\beta}\right) \log\left(\frac{1}{\beta}\right)}\right)$, we have that x appears in Est .

The server uses processing memory of size $\tilde{O}(\sqrt{n})$, and runs in time $\tilde{O}(n)$. Each user has $\tilde{O}(1)$ runtime, memory, and communication.

3.1.2 Non-Private Reduction Based on List Recoverable Codes [17]. Recall that in the reduction of Bassily et al. [2] there is only one hash function h . If that hash function “failed” for some heavy-hitter x (meaning that too many other input elements collide with x), then we cannot identify x and we needed to repeat. Suppose instead that we modify the reduction to introduce an *independent* hash

function h_m for every coordinate $m \in [M]$. The upside is that now, except with exponentially small probability (in M), every heavy-hitter x^* causes at most a small fraction of the hash functions to fail. Then hopefully we are able to detect *most* of the symbols of x^* . If instead of applying this scheme to x^* itself, we apply it to an error-correcting encoding of x^* with constant rate and constant relative distance, then we can indeed recover x^* this way.³ However, aside from the fact that some of the symbols are missing, it now becomes unclear how to concatenate the coordinates of the elements that we are constructing. Specifically, in the construction of Bassily et al., the fact that $h(x^*) = y^*$ appeared throughout all coordinates is what we used to concatenate the different symbols we identified. If the hash functions are independent across m , then this does not work.

We now describe this strategy in more detail, and describe the solution proposed by Larsen et al. [17] to overcome this issue. Fix a heavy-hitter x^* . For every coordinate $m \in [M]$ where h_m does not fail on x^* , we have that the frequency of $(h_m(x^*), x^*[m])$ “stands out” in comparison to every $(h_m(x^*), w')$ where $w' \neq x^*[m]$, meaning that $(h_m(x^*), x^*[m])$ appears in S_m significantly more times than any other $(h_m(x^*), w')$. Hence, for every coordinate $m \in [M]$ and for every hash value $y \in [Y]$ we identify (at most 1) symbol w s.t. (y, w) “stands out” compared to all other (y, w') in S_m . In other words, for every $m \in [M]$ we construct a list $L_m = \{(y, w)\}$ of “stand outs” in S_m . The guarantee is that for every heavy-hitter x^* , we have that $(h_m(x^*), x^*[m])$ appears in L_m for most coordinates $m \in [M]$. If we could single out all occurrences of $(h_m(x^*), x^*[m])$ in these lists, then we could concatenate the symbols $x^*[m]$ and reconstruct x^* . The question is how to match symbols from different lists. As was observed in [12] that studied the related problem of ℓ_1/ℓ_1 compressed sensing, this can be done using *list-recoverable codes*.

Recall that in a (standard) error correcting code, the decoder receives a noisy codeword (represented by, say, M symbols from $[W]$) and recovers a (noiseless) codeword that agrees with the noisy codeword on at least $(1 - \alpha)$ fraction of the coordinates, assuming that such a codeword exists. An error correcting code thus protects a transmission from any adversarial corruption of an α -fraction of a codeword. A *list-recoverable code* protects a transmission from uncertainty as well as corruption. Instead of receiving a single (possibly corrupted) symbol in each coordinate, the decoder for a list-recoverable code receives a short list of size ℓ of possible symbols. Let $L_1, \dots, L_M$ be such a sequence of lists, where list L_m contains the m^{th} symbol of a noisy codeword. The requirement on the decoder is that it is able to recover all possible codewords x that “agree” with at least $(1 - \alpha)$ fraction of the lists, in the sense that the m^{th} symbol of x appears in L_m . Formally,

DEFINITION 3.4 ([13]). *An (α, ℓ, L) -list recoverable code is a pair of mappings (Enc, Dec) where*

$$\text{Enc} : X \rightarrow [Z]^M, \quad \text{and} \quad \text{Dec} : ([Z]^\ell)^M \rightarrow X^L,$$

³We remark that this presentation is oversimplified. In the analysis of [17], the problem is not only that each x^* causes a small fraction of the hash functions to fail. Rather, the problem is that a small fraction of the coordinates are “bad” (since otherwise it may be that each x is failed by very few hash functions, but every h_m ends up “bad” because they fail for some different x^* s).

such that the following holds. Let $L_1, \dots, L_M \in [Z]^\ell$. Then for every $x \in X$ satisfying $|\{m : \text{Enc}(x)_m \in L_m\}| \geq (1 - \alpha)M$ we have that $x \in \text{Dec}(L_1, \dots, L_M)$.

In our case, the lists $L_m = \{(y, w)\}$ contain the “stand outs” of S_m , with the guarantee that for every heavy-hitter x^* , we have that $(h_m(x^*), x^*[m])$ appears in L_m for most coordinates m . We can therefore, in principle, use a list-recoverable code in order to recover all heavy-hitters.⁴ However, there are currently no explicit codes achieving the parameter settings that are needed for the construction. To overcome this issue, Larsen et al. [17] constructed a *relaxed* variant of list-recoverable codes, and showed that their relaxed variant still suffices for their (non-private) reduction. Intuitively, their relaxed code utilizes the randomness in the first portion of every element in the lists (i.e., the outcome of the hash functions). Formally,

DEFINITION 3.5. An (α, ℓ, L) -unique-list-recoverable code is a pair of mappings (Enc, Dec) where

$$\text{Enc} : X \rightarrow ([Y] \times [Z])^M, \quad \text{and} \quad \text{Dec} : ([Y] \times [Z])^\ell \rightarrow X^L,$$

such that the following holds. Let $L_1, \dots, L_M \in ([Y] \times [Z])^\ell$. Assume that for every $m \in [M]$, if $(y, z), (y', z') \in L_m$ then $y \neq y'$. Then for every $x \in X$ satisfying $|\{m : \text{Enc}(x)_m \in L_m\}| \geq (1 - \alpha)M$ we have that $x \in \text{Dec}(L_1, \dots, L_M)$.

THEOREM 3.6 ([17]). There exist constants $C > 1$ and $0 < \alpha < 1$ such that the following holds. For all constants $M \leq \log |X|$ and $Y, \ell \in \mathbb{N}$, and for every fixed choice of functions $h_1, \dots, h_M : X \rightarrow [Y]$, there is a construction of an (α, ℓ, L) -unique-list-recoverable code $\text{Enc} : X \rightarrow ([Y] \times [Z])^M$ and $\text{Dec} : ([Y] \times [Z])^\ell \rightarrow X^L$, where $L \leq C \cdot \ell$ and $Z \leq (|X|^{1/M} \cdot Y)^C$. Furthermore, there is a mapping $\widetilde{\text{Enc}} : X \rightarrow [Z]^M$ such that for every $x \in X$ we have

$$\text{Enc}(x) = \left((h_1(x), \widetilde{\text{Enc}}(x)_1), \dots, (h_M(x), \widetilde{\text{Enc}}(x)_M) \right).$$

There is a pre-processing algorithm which, given $M, Y, \ell, h_1, \dots, h_M$, computes a description of the functions Enc, Dec in $\text{poly}(M)$ time.⁵ Then, evaluating $\text{Enc}(x)$ takes linear time and space in addition to $O(M)$ invocations of the hash functions, and evaluating $\text{Dec}(L_1, \dots, L_M)$ takes linear space and polynomial time.

We remark that the focus in [17] was on space and runtime, and that in fact, in the non-private literature the meaning of “minimal achievable error” is unclear, as achieving zero error is trivial (it is “only” a question of runtime and memory).

3.2 Our Contribution

In the following sections we construct an LDP heavy-hitters protocol and show that it achieves optimal error in all parameters. Our private algorithm is obtained from the non-private algorithm of [17] with the following modifications.

As [17] showed, by appropriately choosing the parameters of a unique-list-recoverable code, the strategy outlined above identifies

⁴In fact, with this formulation it suffices that the lists L_m contain only the set of heavy items $\{w\}$ found to be heavy in S_m .

⁵The pre-processing algorithm has deterministic $\text{poly}(M)$ time for every M of the form $M = D^i$ for some universal constant D . Otherwise, the algorithm has $\text{poly}(M)$ Las Vegas expected runtime. We ignore this issue for simplicity.

all “extra heavy” elements, i.e., those with frequencies $\gtrsim n/\log |X|$. In order to identify Δ -heavy elements for some parameter Δ , the first step in their algorithm is to hash input elements into $\tilde{O}(n/\Delta)$ buckets. Informally, one can then show that w.h.p. the weight of every Δ -heavy element increases significantly relative to the size of its bucket, and in particular becomes “extra-heavy” with respect to its bucket. This allows [17] to focus on identifying “extra heavy” elements within each bucket separately.

Under LDP, however, this reduction to the larger error case cannot be applied directly, as it is unclear at which point to apply the hash function. The server cannot apply it since it does not get to see the users’ inputs, and the users cannot reveal their buckets to the server as this would compromise privacy. We overcome this issue by appending the outcome of a similar hash function to the encoding of elements in the unique-list-recoverable code. As we will see in Section 3.4, this allows us to identify all Δ -heavy elements directly in the original database. We obtain the following theorem:

THEOREM 3.7. Let $\beta, \epsilon \leq 1$, and let $S \in X^n$ be a (distributed) database. There exists an ϵ -LDP algorithm that returns a list Est of length $\tilde{O}(\sqrt{n})$ such that with probability $1 - \beta$

$$(1) \forall (x, a) \in \text{Est}, |a - f_S(x)| \leq O\left(\frac{1}{\epsilon} \sqrt{n \log \left(\frac{\min\{n, |X|\}}{\beta}\right)}\right).$$

$$(2) \text{ For every } x \in X \text{ s.t. } f_S(x) \geq O\left(\frac{1}{\epsilon} \sqrt{n \log \left(\frac{|X|}{\beta}\right)}\right), \text{ we have that } x \text{ is included in Est.}$$

The server uses processing memory of size $\tilde{O}(\sqrt{n})$, and runs in time $\tilde{O}(n)$. Each user has $\tilde{O}(1)$ runtime, memory, and communication.

3.3 Additional Preliminaries

We start by presenting additional preliminaries that enable our construction.

3.3.1 Frequency Oracle. Recall that a frequency oracle is a protocol that, after communicating with the users, outputs a *data structure* capable of approximating the frequency of every domain element $x \in X$. Our protocol uses the frequency oracle of [2], named *Hashtogram*, as a subroutine.

THEOREM 3.8 (Hashtogram [2]). Let $S \in X^n$ be a (distributed) database, and let $\beta, \epsilon \leq 1$. There exists an ϵ -LDP algorithm such that the following holds. Fix a domain element $x \in X$ to be given as a query to the algorithm. With probability at least $1 - \beta$, the algorithm answers the query x with $a(x)$ satisfying:

$$|a(x) - f_S(x)| \leq O\left(\frac{1}{\epsilon} \cdot \sqrt{n \log \left(\frac{\min\{n, |X|\}}{\beta}\right)}\right).$$

The server uses processing memory of size $\tilde{O}(\sqrt{n})$, and runs in time $\tilde{O}(1)$ per query (plus $\tilde{O}(n)$ for instantiation). Each user has $\tilde{O}(1)$ runtime, memory, and communication. The $\tilde{O}$ notation hides logarithmic factors in n and $|X|$.

Observe that by a union bound, algorithm *Hashtogram* answers every fixed set of w queries with worst case error $O\left(\frac{1}{\epsilon} \cdot \sqrt{n \log \left(\frac{w \cdot \min\{n, |X|\}}{\beta}\right)}\right)$. The $\min\{n, |X|\}$ factor is known

to be unnecessary for some regimes of n , $|X|$, β , or if one allows a slightly larger runtime or memory. In particular, Bassily et al. gave a (simplified) analysis of algorithm Hashtogram for the case where the domain X is small, with the following guarantees.

THEOREM 3.9 (Hashtogram [2]). *Let $S \in X^n$ be a (distributed) database, and let $\beta, \varepsilon \leq 1$. There exists an ε -LDP algorithm such that the following holds. Fix a domain element $x \in X$ to be given as a query to the algorithm. With probability at least $1 - \beta$, the algorithm answers the query x with $a(x)$ satisfying:*

$$|a(x) - f_S(x)| \leq O\left(\frac{1}{\varepsilon} \cdot \sqrt{n \log\left(\frac{1}{\beta}\right)}\right).$$

The server uses processing memory of size $\tilde{O}(|X|)$, and runs in time $\tilde{O}(1)$ per query (plus $\tilde{O}(|X|^2)$ for instantiation). Each user has $\tilde{O}(1)$ runtime, memory, and communication. The $\tilde{O}$ notation hides logarithmic factors in n and $|X|$.

We will only apply Theorem 3.9 on small domains X (satisfying $|X|^2 = \tilde{O}(n)$). Our eventual analysis in the proof of Theorem 3.7 will make use of both accuracy guarantees of Hashtogram.

3.3.2 The Poisson Approximation. When throwing n balls into R bins, the distribution of the number of balls in a given bin is $\text{Bin}(n, 1/R)$. As the Poisson distribution is the limit distribution of the binomial distribution when n/R is fixed and $n \rightarrow \infty$, the distribution of the number of balls in a given bin is approximately $\text{Pois}(n/R)$. In fact, in some cases we may approximate the *joint distribution* of the number of balls in all the bins by assuming that the load in each bin is an independent Poisson random variable with mean n/R .

THEOREM 3.10 (E.G., [19]). *Suppose that n balls are thrown into R bins independently and uniformly at random, and let X_i be the number of balls in the i^{th} bin, where $1 \leq i \leq R$. Let $Y_1, \dots, Y_R$ be independent Poisson random variables with mean n/R . Let $f(x_1, \dots, x_R)$ be a nonnegative function. Then,*

$$\mathbb{E}[f(X_1, \dots, X_R)] \leq e\sqrt{n} \mathbb{E}[f(Y_1, \dots, Y_R)].$$

In particular, the theorem states that any event that takes place with probability p in the Poisson case, takes place with probability at most $pe\sqrt{n}$ in the exact case (this follows by letting f be the indicator function of that event).

We will also use the following bounds for the tail probabilities of a Poisson random variable:

THEOREM 3.11 ([1]). *Let X have Poisson distribution with mean μ . For $0 \leq \alpha \leq 1$,*

$$\begin{aligned} \Pr[X \leq \mu(1 - \alpha)] &\leq e^{-\alpha^2 \mu/2} \\ \Pr[X \geq \mu(1 + \alpha)] &\leq e^{-\alpha^2 \mu/2}. \end{aligned}$$

3.3.3 Concentration (with Limited Independence). We first state a version of the multiplicative Chernoff bound, including a formulation of the upper tail bound which holds under limited independence.

THEOREM 3.12 ([23, THEOREM 2]). *Let $X_1, \dots, X_n$ be random variables taking values in $\{0, 1\}$. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$. Then for every $0 \leq \alpha \leq 1$,*

(1) *If the X_i 's are $\lceil \mu\alpha \rceil$ -wise independent, then*

$$\Pr[X \geq \mu(1 + \alpha)] \leq \exp(-\alpha^2 \mu/3).$$

(2) *If the X_i 's are fully independent, then*

$$\Pr[X \leq \mu(1 - \alpha)] \leq \exp(-\alpha^2 \mu/2).$$

Next, we present a limited independence version of Bernstein's inequality.

THEOREM 3.13 ([15, LEMMA 2]). *Let $X_1, \dots, X_n$ be k -wise independent random variables, for an even integer $k \geq 2$, that are each bounded by T in magnitude. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$. Let $\sigma^2 = \sum_i \mathbb{E}(X_i - \mathbb{E}X_i)^2$. Then there exists a constant $C > 0$ such that for all $\lambda > 0$,*

$$\Pr[|X - \mu| > \lambda] \leq C^k \cdot \left((\sigma\sqrt{k}/\lambda)^k + (Tk/\lambda)^k \right).$$

3.4 The Full Protocol

In this section we construct and analyze our heavy-hitters algorithm, and prove Theorem 3.7. The full algorithm appears as Algorithm PrivateExpanderSketch.

It is immediate from the construction that algorithm PrivateExpanderSketch satisfies ε -LDP. We now proceed with the utility analysis. While we must make the restriction $n \leq |X|$ for the analysis to go through, we remark that in the complementary case $n > |X|$, it is possible to just apply the frequency oracle Hashtogram (Theorem 3.8) to every item in X and obtain the same runtime, memory, and communication guarantees.

OF THEOREM 3.7. Item 1 of the theorem follows directly from Theorem 3.8 (the properties of Hashtogram). We now prove item 2. Below, we will freely assume that $n \geq O(\log |X|/\varepsilon^2)$, with item 2 being vacuously true otherwise. Consider the execution of PrivateExpanderSketch on a (distributed) database S . We begin by defining a sequence of events $E_1, \dots, E_7$. We will show that all of these events occur simultaneously with high probability, and that if they indeed all occur, then the guarantee of item 2 holds.

For some constant $C_H \geq 1$, we say that an element $x \in X$ is *heavy* if $f_S(x) \geq \frac{C_H}{\varepsilon} \sqrt{n} \cdot \log |X|$, and *non-heavy* otherwise. We let H denote the set of all heavy elements. Observe that $|H| \leq \varepsilon \cdot \sqrt{n/\log |X|}$. For $b \in [B]$ denote the set of all heavy elements that are mapped to b under g as $H_b = \{x \in H : g(x) = b\}$. Consider the following event:

Event E_1 (over the choice of g):

For every $b \in [B]$, the number of heavy elements that are mapped to b under g is at most $C_1 \log |X|$. Furthermore, the sum of frequencies of all non-heavy elements (i.e., elements $x \notin H$) that are mapped to b is at most $\frac{C_1}{\varepsilon} \sqrt{n} \log^{3/2} |X|$, where $C_1 > 0$ is an absolute constant. That is, $\forall b \in [B]$:

$$|H_b| \leq C_1 \log |X| \quad \text{and} \quad \sum_{x \notin H: g(x)=b} f_S(x) \leq \frac{C_1 \sqrt{n} \log^{3/2} |X|}{\varepsilon}$$

We now show that E_1 happens with high probability. For $B = \left\lceil \frac{1}{10C_H} \cdot \varepsilon \sqrt{n/\log^{3/2} |X|} \right\rceil \leq |X|$, it follows by the limited independence Chernoff bound (Theorem 3.12, item 1) and a union bound

Algorithm 1 PrivateExpanderSketch

Notation: Let C_M, C_Y, C_ℓ, C_g , and C_f be universal constants (to be specified later).

Denote $M = C_M \cdot \log |X| / \log \log |X|$, $Y = \log^{C_Y} |X|$, and $\ell = C_\ell \cdot \log |X|$.

Public randomness: Random partition of $[n]$ into M sets $I_1, \dots, I_M$.

Pairwise independent hash functions $h_1, \dots, h_M : X \rightarrow [Y]$.

$(C_g \cdot \log |X|)$ -wise independent hash function $g : X \rightarrow [B]$.

Tool used: A unique-list-recoverable code (Enc, Dec) with parameters M, Y, ℓ using $h_1, \dots, h_M$, as in Theorem 3.6.

Setting: Each player $i \in [n]$ holds a value $x_i \in X$. Define $S = (x_1, \dots, x_n)$.

For $m \in [M]$, let $S_m = \left(g(x_i), \text{Enc}(x_i)_m \right)_{i \in I_m} = \left(g(x_i), h_m(x_i), \widehat{\text{Enc}}(x_i)_m \right)_{i \in I_m}$.

1. $\forall m \in [M]$, use Hashtogram(S_m) with privacy parameter $\frac{\epsilon}{2}$ to get estimates $\hat{f}_{S_m}(b, y, z)$ for $f_{S_m}(b, y, z)$ for every $(b, y, z) \in [B] \times [Y] \times [Z]$.
2. For every $(m, b) \in [M] \times [B]$, initialize $L_m^b = \emptyset$.
3. For every $(m, b, y) \in [M] \times [B] \times [Y]$:
 - (a) Let $z = \arg\max \{ \hat{f}_{S_m}(b, y, z) \}$.
 - (b) If $\hat{f}_{S_m}(b, y, z) \geq C_f \cdot \frac{\log \log |X|}{\epsilon} \cdot \sqrt{\frac{n}{\log |X|}} = O\left(\frac{\sqrt{n \log |X|}}{\epsilon M}\right)$ and $|L_m^b| \leq C_\ell \cdot \log |X|$ then add (y, z) to L_m^b .
4. For every $b \in [B]$ let $\widehat{H}^b = \text{Dec}(L_1^b, \dots, L_M^b)$. Define $\widehat{H} = \bigcup_{b \in [B]} \widehat{H}^b$.
5. Use Hashtogram(S) with privacy parameter $\frac{\epsilon}{2}$ to obtain $\hat{f}_S(x) \approx f_S(x)$ for every $x \in \widehat{H}$.
6. Return $\text{Est} = \left\{ (x, \hat{f}_S(x)) : x \in \widehat{H} \right\}$.

over all $b \in [B]$ that for some constant $C_1 > 0$,

$$\Pr_g \left[\begin{array}{l} \exists b \in [B] \text{ s.t.} \\ |\{x \in H : g(x) = b\}| > C_1 \log |X| \end{array} \right] < \frac{1}{2|X|}. \quad (1)$$

Here, we made use of the fact that $g : X \rightarrow [B]$ is $(C_g \cdot \log |X|)$ -wise independent for a constant C_g . We next invoke the limited independence formulation of Bernstein's inequality (Theorem 3.13). Fix $b \in [B]$. We will consider the collection of random variables R_x indexed by $x \in X \setminus H$ defined as follows. Define $R_x = \mathbf{1}_{g(x)=b} \cdot f_S(x)$. Then each R_x is bounded by $T < \frac{1}{\epsilon} \sqrt{n \cdot \log |X|}$ and $\sum_{x \notin H} R_x$ has variance

$$\begin{aligned} \sigma^2 &= \sum_{x \notin H} f_S^2(x) \left(\frac{1}{B} - \frac{1}{B^2} \right) \leq \sum_{x \notin H} \frac{1}{B} \cdot f_S^2(x) \\ &\leq \frac{1}{\epsilon^2} n \log^2 |X|. \end{aligned}$$

It then follows by Bernstein's inequality and a union bound over all $b \in [B]$ that

$$\Pr_g \left[\begin{array}{l} \exists b \in [B] \text{ s.t.} \\ \sum_{x \notin H: g(x)=b} f_S(x) > \frac{C_1}{\epsilon} \sqrt{n \log^3 |X|} \end{array} \right] < \frac{1}{2|X|}.$$

This shows that $\Pr[E_1] \geq 1 - \frac{1}{|X|}$. We now continue the analysis assuming that Event E_1 occurs.

Event E_2 (over the choice of $I_1, \dots, I_M$):

For every $(b, m) \in [B] \times [M]$ we have

$$|\{i \in I_m : x_i \notin H \text{ and } g(x_i)=b\}| \leq \frac{2C_1}{M\epsilon} \sqrt{n \log^3 |X|}.$$

Conditioned on Event E_1 , for every $b \in [B]$ we have that $|\{i \in [n] : x_i \notin H \text{ and } g(x_i) = b\}| \leq \frac{C_1}{\epsilon} \sqrt{n \log^3 |X|}$. Hence, by the Chernoff bound (Theorem 3.12, item 2) and a union bound over all $(b, m) \in [B] \times [M]$, we get that $\Pr[E_2 | E_1] \geq 1 - \frac{1}{|X|}$.

Event E_3 (over the choice of $I_1, \dots, I_M$):

For every $x \in H$ there exists a subset $M_3^x \subseteq [M]$ of size $|M_3^x| \geq (1 - \frac{\alpha}{10})M$ s.t. for every $m \in M_3^x$ we have that $|\{i \in I_m : x_i = x\}| \geq \frac{f_S(x)}{2M}$.

Recall that $M = C_M \cdot \log |X| / \log \log |X|$. We now analyze event E_3 using the Poisson approximation for balls into bins (see Theorem 3.10). To that end, fix $x \in H$, and let $\tilde{J}_1, \dots, \tilde{J}_M$ be independent Poisson random variables with mean $\frac{f_S(x)}{M}$. Let us say that $m \in [M]$ is *bad* if $\tilde{J}_m < \frac{f_S(x)}{2M}$. Now fix $m \in [M]$. Using a tail bound for the Poisson distribution (see Theorem 3.11), assuming that $n \geq 8C_M^2 \log |X|$, we have that m is bad with probability at most $\frac{1}{\log^{C_M} |X|}$. As $\tilde{J}_1, \dots, \tilde{J}_M$ are independent, the probability that there are more than $\alpha M / 10$ bad choices for m is at most

$$\begin{aligned} \binom{M}{\alpha M / 10} \left(\frac{1}{\log^{C_M} |X|} \right)^{\alpha M / 10} &\leq \left(\frac{10e}{\alpha \log^{C_M} |X|} \right)^{\frac{\alpha C_M \log |X|}{10 \log \log |X|}} \\ &= \left(\frac{10e}{\alpha} \right)^{\frac{\alpha C_M \log |X|}{10 \log \log |X|}} \cdot \left(\frac{1}{|X|} \right)^{\frac{\alpha C_M^2}{10}} \\ &= \left(\frac{1}{|X|} \right)^{\frac{\alpha C_M^2}{10} - o(1)}. \end{aligned}$$

Hence, by the Poisson approximation (Theorem 3.10) we get that

$$\Pr \left[\begin{array}{l} \text{There are at least } \alpha M/10 \text{ indices } m \text{ s.t.} \\ |\{i \in I_m : x_i = x\}| < f_S(x)/(2M) \end{array} \right] \\ \leq e\sqrt{f_S(x)} \cdot \left(\frac{1}{|X|} \right)^{\frac{\alpha C_M^2}{10} - o(1)}.$$

Choosing $C_M \geq \frac{10}{\alpha}$ and using the union bound over every $x \in H$ we get that $\Pr[E_3] \geq 1 - \frac{1}{|X|}$.

Event E_4 (over the choice of $I_1, \dots, I_M$):

There exists a subset $M_4 \subseteq [M]$ of size $|M_4| \geq (1 - \frac{\alpha}{10})M$ s.t. for every $m \in M_4$ we have

$$\frac{n}{2M} \leq |I_m| \leq \frac{2n}{M}.$$

An analysis similar to that of Event E_3 shows that $\Pr[E_4] \geq 1 - \frac{1}{|X|}$.

Event E_5 (over sampling $h_1, \dots, h_M$):

For every $b \in [B]$ there exists a subset $M_5^b \subseteq [M]$ of size $|M_5^b| \geq (1 - \frac{\alpha}{10})M$ s.t. for every $m \in M_5^b$ we have that h_m perfectly hashes the elements of $H_b = \{x \in H : g(x) = b\}$.

We analyze event E_5 assuming that event E_1 occurs, in which case $|H_b| \leq C_1 \cdot \log |X|$ for every $b \in [B]$. Fix $b \in [B]$ and $m \in [M]$. Recall that h_m is a pairwise independent hash function mapping X to $[Y]$. Hence, for $x' \neq x$ we have $\Pr_{h_m}[h_m(x') = h_m(x)] = \frac{1}{Y}$. Using the union bound over every pair (x, x') with $x' \neq x \in H_b$, assuming that $Y \geq C_1^2 \cdot \log^{C_M+2} |X|$, we have

$$\Pr_{h_m}[\exists(x, x') \in H_b \times H_b \text{ s.t. } x \neq x', h_m(x') = h_m(x)] \\ \leq \frac{1}{Y} \binom{|H_b|}{2} \leq \frac{1}{\log^{C_M} |X|}.$$

That is, the probability that h_m does not perfectly hash the elements of H_b is at most $\frac{1}{\log^{C_M} |X|}$. As the hash functions $h_1, \dots, h_M$ are independent, the probability that there are more than $\alpha M/10$ choices of m s.t. h_m does not perfectly hash H_b is at most

$$\binom{M}{\alpha M/10} \left(\frac{1}{\log^{C_M} |X|} \right)^{\alpha M/10} \leq \left(\frac{1}{|X|} \right)^{\frac{\alpha C_M^2}{10} - o(1)}.$$

Setting $C_M \geq \frac{10}{\alpha}$ and using the union bound over every $b \in [B]$ we get that $\Pr[E_5|E_1] \geq 1 - \frac{1}{|X|}$.

Event E_6 (over sampling $h_1, \dots, h_M$):

For every $x \in H$ there exists a subset $M_6^x \subseteq [M]$ of size $|M_6^x| \geq (1 - \frac{\alpha}{10})M$ s.t. for every $m \in M_6^x$ we have that

$$\left| \left\{ i \in I_m : \begin{array}{l} x_i \notin H, \\ g(x_i) = g(x), \\ h_m(x_i) = h_m(x) \end{array} \right\} \right| \leq \frac{\sqrt{n \cdot \log |X|}}{\epsilon M}.$$

We analyze event E_6 assuming that event E_2 occurs, in which case for every b, m we have that

$$|\{i \in I_m : x_i \notin H, g(x_i) = b\}| \leq \frac{2C_1}{M\epsilon} \sqrt{n \log^{3/2} |X|}.$$

Fix $x \in H$ and fix $m \in [M]$. We have that

$$\mathbb{E}_{h_m} [|\{i \in I_m : x_i \notin H, g(x_i) = g(x), h_m(x_i) = h_m(x)\}|] \\ = \sum_{\substack{i \in I_m: \\ x_i \notin H, \\ g(x_i) = g(x)}} \mathbb{E}_{h_m} [\mathbb{1}_{h_m(x_i) = h_m(x)}] \\ \leq \frac{1}{Y} \cdot \frac{2C_1}{M\epsilon} \sqrt{n \log^{3/2} |X|}.$$

Thus, by Markov's inequality, we have that

$$\Pr_{h_m} \left[\left| \left\{ i \in I_m : \begin{array}{l} x_i \notin H, \\ g(x_i) = g(x), \\ h_m(x_i) = h_m(x) \end{array} \right\} \right| \geq \frac{\sqrt{n \cdot \log |X|}}{\epsilon M} \right] \leq \frac{2C_1 \log |X|}{Y}.$$

Let us say that m is *bad* for x if

$$\left| \left\{ i \in I_m : \begin{array}{l} x_i \notin H, \\ g(x_i) = g(x), \\ h_m(x_i) = h_m(x) \end{array} \right\} \right| \geq \frac{\sqrt{n \cdot \log |X|}}{\epsilon M}.$$

So m is bad with probability at most $\frac{2C_1 \log |X|}{Y}$. As the hash functions $h_1, \dots, h_M$ are independent, assuming that $Y \geq 2C_1 \cdot \log^{C_M+1} |X|$, the probability that there are more than $\alpha M/10$ bad choices for m is at most

$$\binom{M}{\alpha M/10} \left(\frac{1}{\log^{C_M} |X|} \right)^{\alpha M/10} \leq \left(\frac{1}{|X|} \right)^{\frac{\alpha C_M^2}{10} - o(1)}.$$

Setting $C_M \geq \frac{10}{\alpha}$ and using the union bound over every $x \in H$ we get that $\Pr[E_6|E_2] \geq 1 - \frac{1}{|X|}$.

Event E_7 (over randomness in Hashtogram):

For every $b \in [B]$ there exists a subset $M_7^b \subseteq [M]$ of size $|M_7^b| \geq (1 - \frac{\alpha}{5})M$ s.t. for every $m \in M_7^b$ and every $(y, z) \in [Y] \times [Z]$ we have that $|\hat{f}_{S_m}(b, y, z) - f_{S_m}(b, y, z)| \leq \frac{C_7 \cdot \log \log |X|}{\epsilon} \sqrt{n / \log |X|}$, where $C_7 > 0$ is an absolute constant.

We analyze Event E_7 conditioned on Event E_4 , and let $M_4 \subseteq [M]$ denote the set from Event E_4 . Recall that $|M_4| \geq (1 - \frac{\alpha}{10})M$, and that for every $m \in M_4$ we have that $n/2M \leq |S_m| \leq 2n/M$.

Fix $b \in [B]$ and fix $m \in M_4$. Let us say that m is *bad* if there exists $(y, z) \in [Y] \times [Z]$ such that

$$|\hat{f}_{S_m}(b, y, z) - f_{S_m}(b, y, z)| > \frac{6}{\epsilon} \sqrt{|S_m| \cdot \log(4YZ \cdot \log^{C_M} |X|)} \\ = O \left(\frac{\log \log |X|}{\epsilon} \sqrt{n / \log |X|} \right).$$

By the properties of algorithm Hashtogram (Theorem 3.9), we have that m is bad with probability at most $\frac{1}{\log^{C_M} |X|}$. As the coins of Hashtogram are independent across different executions (i.e., for

different values of m), the probability that there are more than $\alpha|M_4|/10$ bad choices for $m \in M_4$ is at most

$$\binom{|M_4|}{\alpha|M_4|/10} \left(\frac{1}{\log^{C_M} |X|} \right)^{\alpha|M_4|/10} \leq \left(\frac{1}{|X|} \right)^{\frac{\alpha C_M^2}{20} - o(1)}.$$

Setting $C_M \geq \frac{20}{\alpha}$ and using the union bound over every $b \in [B]$ we get that $\Pr[E_7|E_4] \geq 1 - \frac{1}{|X|}$.

We are now ready to complete the proof, by showing that whenever $E_1, E_2, E_3, E_4, E_5, E_6, E_7$ occur we have that every heavy element $x \in H$ appears in the output list $\hat{H}$. To that end, fix $x \in H$, and denote $M^x = M_3^x \cap M_4 \cap M_5^{g(x)} \cap M_6^x \cap M_7^{g(x)}$. Observe that $|M^x| \geq (1-\alpha)M$. We now show that for every $m \in M^x$ it holds that $(h_m(x), \widetilde{\text{Enc}}(x)_m) \in L_m^{g(x)}$, in which case, by the properties of the code (Enc, Dec) we have that $x \in \hat{H}^{g(x)} = \text{Dec}(L_1^{g(x)}, \dots, L_M^{g(x)})$, and hence, $x \in \hat{H}$.

So let $m \in M^x$. By Event E_3 we have that

$$f_{S_m}(g(x), h_m(x), \widetilde{\text{Enc}}(x)_m) \geq \frac{C_H \cdot \log \log |X|}{2C_M \cdot \epsilon} \sqrt{n/\log |X|}.$$

Hence, by Event E_7 we have that

$$\hat{f}_{S_m}(g(x), h_m(x), \widetilde{\text{Enc}}(x)_m) \geq C_f \cdot \frac{\log \log |X|}{\epsilon} \sqrt{n/\log |X|}, \quad (2)$$

where $C_f = \left(\frac{C_H}{2C_M} - C_7 \right)$.

On the other hand, let $z \neq \widetilde{\text{Enc}}(x)_m$, and recall that by Event E_5 there does not exist a heavy element $x' \neq x$ s.t. $g(x') = g(x)$ and $h_m(x') = h_m(x)$. In addition, by Event E_6 , there could be at most $\frac{\log \log |X|}{\epsilon} \sqrt{n/\log |X|}$ indices $i \in I_m$ s.t. x_i is non-heavy and $g(x_i) = g(x)$ and $h_m(x_i) = h_m(x)$. That is,

$$\begin{aligned} f_{S_m}(g(x), h_m(x), z) &= \left| \left\{ i \in I_m : \begin{array}{l} x_i \in H, \\ g(x_i) = g(x), \\ h_m(x_i) = h_m(x), \\ \widetilde{\text{Enc}}(x_i)_m = z \end{array} \right\} \right| \\ &+ \left| \left\{ i \in I_m : \begin{array}{l} x_i \notin H, \\ g(x_i) = g(x), \\ h_m(x_i) = h_m(x), \\ \widetilde{\text{Enc}}(x_i)_m = z \end{array} \right\} \right| \\ &\leq 0 + \frac{\log \log |X|}{\epsilon} \sqrt{n/\log |X|}. \end{aligned}$$

Hence, by Event E_7 we have that

$$\hat{f}_{S_m}(g(x), h_m(x), z) \leq (C_7 + 1) \frac{\log \log |X|}{\epsilon} \sqrt{n/\log |X|}. \quad (3)$$

By Inequalities (2) and (3), and by setting C_H large enough, we ensure that $\widetilde{\text{Enc}}(x)_m$ is identified in step 3a as the $\arg\max$ of $\hat{f}_{S_m}(g(x), h_m(x), \cdot)$. Hence, in order to show that $(h_m(x), \widetilde{\text{Enc}}(x)_m)$ is added to $L_m^{g(x)}$ in step 3b, it suffices to show that $|L_m^{g(x)}| \leq C_\ell \cdot \log |X|$ for some constant C_ℓ . First recall that, by Event E_1 , there are at most $C_1 \cdot \log |X|$ heavy elements $x' \in H$ s.t. $g(x') = g(x)$. Hence, there could be at most $C_1 \cdot \log |X|$ pairs $(y, z) \in L_m^{g(x)}$ s.t. $\exists x' \in H$ s.t.

$$(g(x'), h_m(x'), \widetilde{\text{Enc}}_m(x')) = (g(x), y, z).$$

Therefore, it suffices to show that $L_m^{g(x)}$ contains at most $O(\log |X|)$ pairs (y, z) that are added because they do *not* correspond to a heavy element $x' \in H$.

To that end, observe that by the condition on step 3b, and by Event E_7 , in order for a pair (y, z) to be added to $L_m^{g(x)}$ it must be that

$$f_{S_m}(g(x), y, z) \geq \frac{\log \log |X|}{\epsilon} \sqrt{n/\log |X|}.$$

However, by Event E_2 , there are at most

$$\frac{2C_1 \cdot \log \log |X|}{C_M \cdot \epsilon} \sqrt{n \log |X|}$$

elements in S_m that yield the same value $g(x)$, but are not generated from a heavy element. Hence, there could be at most $\frac{2C_1}{C_M} \cdot \log |X|$ pairs $(y, z) \in L_m^{g(x)}$ such that there is no $x' \in H$ where $(g(x'), h_m(x'), \widetilde{\text{Enc}}_m(x')) = (g(x), y, z)$. Overall we have that $|L_m^{g(x)}| \leq C_\ell \cdot \log |X|$ for $C_\ell = C_1 + 2C_1/C_M$. $\square$

4 ADVANCED GROUPOSITION AND MAX-INFORMATION

In this section, we show that local differential privacy admits stronger guarantees of group privacy than differential privacy in the central model. In particular, the protection offered to groups of size k under LDP degrades proportionally to about $\sqrt{k}$, a quadratic improvement over what happens in the central model. It is not a coincidence that this behavior mirrors the privacy guarantee of differentially private algorithms under composition. Indeed, the proof of our “advanced” group privacy bound will follow the same strategy as the proof of the so-called advanced composition theorem [9, 10].

DEFINITION 4.1. For random variables A, B , define the privacy loss function

$$\ell_{A,B}(y) = \ln(\Pr[A = y]/\Pr[B = y]).$$

Define the privacy loss random variable $L_{A,B}$ to be $\ell_{A,B}(y)$ for $y \leftarrow A$.

Just as in the proof of the advanced composition theorem, we will leverage the fact that the expected privacy loss of any individual local randomizer is $L_{R_i(x_i), R_i(x_i)} = O(\epsilon^2)$, which is substantially smaller than the worst-case privacy loss of ϵ . Since each local randomizer is applied independently, the cumulative privacy loss incurred by changing a group of k inputs concentrates to within $O(\sqrt{k}\epsilon)$ of its expectation $O(k\epsilon^2)$.

THEOREM 4.2 (ADVANCED GROUPOSITION FOR PURE LDP). Let $x \in X^n$ and $x' \in X^n$ differ in at most k entries for some $1 \leq k \leq n$. Let $\mathcal{A} = (R_1, \dots, R_n) : X^n \rightarrow Y$ be ϵ -LDP. Then for every $\delta > 0$ and $\epsilon' = k\epsilon^2/2 + \epsilon\sqrt{2k \ln(1/\delta)}$, we have

$$\Pr_{y \leftarrow \mathcal{A}(x)} \left[\ln \frac{\Pr[\mathcal{A}(x) = y]}{\Pr[\mathcal{A}(x') = y]} > \epsilon' \right] \leq \delta.$$

In particular, for every $\delta > 0$ and every set $T \subseteq Y$, we have $\Pr[\mathcal{A}(x) \in T] \leq e^{\epsilon'} \Pr[\mathcal{A}(x') \in T] + \delta$.

PROOF. Without loss of generality, we may assume that x and x' differ in the first k coordinates. Since each randomizer is applied

independently, we may write the privacy loss between $\mathcal{A}(x)$ and $\mathcal{A}(x')$ as

$$L_{\mathcal{A}(x), \mathcal{A}(x')} = \sum_{i=1}^k L_{R_i(x_i), R_i(x'_i)}$$

Since each R_i is ε -differentially private, Proposition 3.3 of [4] implies

$$\mathbb{E} \left[L_{R_i(x_i), R_i(x'_i)} \right] \leq \frac{1}{2} \varepsilon^2, \quad |L_{R_i(x_i), R_i(x'_i)}| \leq \varepsilon.$$

Hence, by Hoeffding's inequality, we have that for every $t > 0$,

$$\Pr[L_{R(x), R(x')} > k\varepsilon^2/2 + t] \leq e^{-t^2/2k\varepsilon^2}.$$

Setting $t = \varepsilon\sqrt{2k \ln(1/\delta)}$ completes the proof. $\square$

We remark that it is straightforward to extend Theorem 4.2 to handle (ε, δ) -LDP algorithms using the same ideas as in the proof of the advanced composition theorem (see the discussion in Section 2.2 of [26]).

THEOREM 4.3 (ADVANCED GROUP POSITION FOR APPROXIMATE LDP). *Let $x \in X^n$ and $x' \in X^n$ differ in at most k entries for some $1 \leq k \leq n$. Let $\mathcal{A} = (R_1, \dots, R_n) : X^n \rightarrow Y$ be (ε, δ) -LDP. Then for every $\delta' > 0$, $\varepsilon' = k\varepsilon^2/2 + \varepsilon\sqrt{2k \ln(1/\delta')}$, and every set $T \subseteq Y$, we have $\Pr[\mathcal{A}(x) \in T] \leq e^{\varepsilon'} \Pr[\mathcal{A}(x') \in T] + \delta + k\delta'$.*

Our improved group privacy bound immediately implies a strong bound on the *max-information* of an LDP protocol. The max information [7] of an algorithm is a measure of how much information it reveals about a randomly chosen input. The motivation for studying max-information comes from its ability to ensure generalization in adaptive data analysis.

DEFINITION 4.4. *Let Z, W be jointly distributed random variables. We say that the β -approximate max-information between Z and W , denoted $I_{\infty}^{\beta}(Z; W)$, is at most k if for every measurable subset T of the support of (Z, W) with $\Pr[(Z, W) \in T] > \beta$, we have*

$$\ln \frac{\Pr[(Z, W) \in T] - \beta}{\Pr[Z \otimes W \in T]} \leq k.$$

Here, $Z \otimes W$ denotes the product distribution formed from the marginals Z and W . Given an algorithm $\mathcal{A} : X^n \rightarrow \mathcal{R}$, we say the β -approximate max-information of $\mathcal{A}$, denoted $I_{\infty}^{\beta}(\mathcal{A}, n) \leq k$ if $I_{\infty}^{\beta}(\mathcal{D}; \mathcal{A}(\mathcal{D})) \leq k$ for every distribution $\mathcal{D}$ on X^n .

Dwork et al. [7] showed that ε -DP algorithms have max information $O(\varepsilon n)$. Moreover, they have β approximate max-information $O(\varepsilon^2 n + \varepsilon\sqrt{n \ln(1/\beta)})$, but only when the input distribution $\mathcal{D}$ is a product distribution. Subsequent work by Rogers et al. [22] extended this analysis to (ε, δ) -DP algorithms, for which they showed the restriction to product distributions to be necessary.

Using our group privacy bound for local differential privacy, we show that even under non-product distributions, the max information of ε -LDP protocols has the same behavior as ε -DP algorithms on product distributions. This provides a distinct advantage in adaptive data analysis, as it means ε -LDP algorithms can be composed with arbitrary low max-information algorithms *in any order* while giving strong generalization guarantees.

THEOREM 4.5. *Let $\mathcal{A} : X^n \rightarrow Y$ be ε -LDP. Then for every $\beta > 0$, we have $I_{\infty}^{\beta}(\mathcal{A}, n) \leq n\varepsilon^2/2 + \varepsilon\sqrt{2n \ln(1/\beta)}$.*

PROOF. Let $\mathcal{D}$ be a distribution over X^n . Then for every $k > 0$, by Jensen's inequality,

$$\begin{aligned} & \Pr_{x \sim \mathcal{D}, y \leftarrow \mathcal{A}(x)} \left[\ln \frac{\Pr[\mathcal{D} = x, \mathcal{A}(x) = y]}{\Pr[\mathcal{D} = x] \Pr[\mathcal{A}(\mathcal{D}) = y]} > k \right] \\ &= \Pr_{x \sim \mathcal{D}, y \leftarrow \mathcal{A}(x)} \left[\ln \frac{\Pr[\mathcal{A}(x) = y]}{\Pr[\mathcal{A}(\mathcal{D}) = y]} > k \right] \\ &= \Pr_{x \sim \mathcal{D}, y \leftarrow \mathcal{A}(x)} \left[\ln \frac{\Pr[\mathcal{A}(x) = y]}{\mathbb{E}_{x' \sim \mathcal{D}} [\Pr[\mathcal{A}(x') = y]]} > k \right] \\ &\leq \Pr_{x \sim \mathcal{D}, y \leftarrow \mathcal{A}(x)} \left[\mathbb{E}_{x' \sim \mathcal{D}} \left[\ln \frac{\Pr[\mathcal{A}(x) = y]}{\Pr[\mathcal{A}(x') = y]} \right] > k \right] \\ &\leq \max_{x, x'} \Pr_{y \leftarrow \mathcal{A}(x)} \left[\ln \frac{\Pr[\mathcal{A}(x) = y]}{\Pr[\mathcal{A}(x') = y]} > k \right]. \end{aligned}$$

By Theorem 4.2, this quantity is at most β for $k = n\varepsilon^2/2 + \varepsilon\sqrt{2n \ln(1/\beta)}$.

The claim now follows from the general fact that, for any pair of random variables U, V over the same sample space, $\Pr_{u \sim U} \left[\ln \frac{\Pr[U=u]}{\Pr[V=u]} > k \right] \leq \beta$ implies that

$$\sup_{T: \Pr[U \in T] > \beta} \ln \frac{\Pr[U \in T] - \beta}{\Pr[V \in T]} \leq k.$$

$\square$

5 COMPOSITION FOR RANDOMIZED RESPONSE

In this section, we give a direct proof showing that k instantiations of randomized response obey privacy guarantees matching advanced composition, even under pure ε -LDP. Specifically, we exhibit for every $\beta > 0$ an $O(\varepsilon\sqrt{k \ln(1/\beta)})$ -LDP algorithm that when run on any given input, yields a distribution which is β -close in statistical distance to the composition of k instances of randomized response.⁶ Prior work of Duchi, Jordan, and Wainwright [6] showed that the *problem* of estimating k binary attributes admits an LDP algorithm with similar guarantees as ours. However, their distribution does not resemble the k -fold composition of randomized response. We interpret our new construction as evidence that an advanced composition theorem for pure ε -LDP might hold for more general, and potentially even interactive mechanisms.

Let $M_i : \{0, 1\}^k \rightarrow \{0, 1\}$ that performs randomized response on the i th bit of its input. On an input x , our algorithm first identifies a “good” set G_x of outputs $(y_1, \dots, y_k) \in \{0, 1\}^k$ which have roughly average likelihood of appearing under the composition $M(x) = (M_1(x), \dots, M_k(x))$. By concentration arguments, $M(x)$ produces an outcome in G_x with probability at least $1 - \beta$. Our approximate composed algorithm $\tilde{M}(x)$ simply runs $M(x)$, returns the outcome if it is in G_x , and returns a uniformly random element outside G_x otherwise. The crux of the privacy argument is to show that elements in this latter case end up being sampled with roughly the same probability as elements in G_x .

We obtain the following result. See Algorithm 2 for the construction. The analysis is omitted from this extended abstract.

⁶The usual guarantee of advanced composition also includes an additive $O(\varepsilon^2 k)$ term, but our result only applies to the typical setting where $\varepsilon^2 k \leq \varepsilon\sqrt{k}$.

Algorithm 2 Approximate composed algorithm $\tilde{M}(x)$

-
1. Sample $y = (y_1, \dots, y_k) \leftarrow M(x)$
 2. If $y \in G_x$, output y ; Else, output a random sample from $\{0, 1\}^k \setminus G_x$.
-

THEOREM 5.1. For each $i \in [k]$, let $M_i : \{0, 1\}^k \rightarrow \{0, 1\}$ be the ϵ -differentially private algorithm

$$M_i(x) = \begin{cases} x_i & \text{w.p. } \frac{e^\epsilon}{e^\epsilon + 1} \\ 1 - x_i & \text{w.p. } \frac{1}{e^\epsilon + 1}. \end{cases}$$

Let $0 < \beta < \left(\frac{\epsilon\sqrt{k}}{2(k+1)}\right)^{2/3}$ and suppose $\tilde{\epsilon} = 6\epsilon\sqrt{k \ln(\frac{1}{\beta})} \leq 1$. Then there exists an algorithm $\tilde{M} : \{0, 1\}^k \rightarrow \{0, 1\}^k$ such that

- (1) $\tilde{M}$ is $\tilde{\epsilon}$ -differentially private.
- (2) For every $x \in \{0, 1\}^k$, there exists an event E with $\Pr[E] \geq 1 - \beta$ such that, conditioned on E , the output $\tilde{M}(x)$ is identically distributed to $M(x) = (M_1(x), \dots, M_k(x))$.

6 GENERIC TRANSFORMATION FROM APPROXIMATE TO PURE PRIVACY

In this section we present our generic transformation from any (non-interactive) (ϵ, δ) -LDP protocol into an $O(\epsilon)$ -LDP protocol with essentially the same utility guarantees. Our transformation is based on *rejection sampling*, an idea that first appeared in the context of differential privacy in the work of Kasiviswanathan et al. [16], where it was used to show an equivalence between the local model (under pure-privacy) and SQ learning. Rejection sampling was also used by Bassily and Smith [3] to show that almost any (non-interactive) ϵ -LDP protocol can be transformed into an ϵ -LDP protocol where the communication per user is only 1 bit (at the expense of increasing the shared randomness in the protocol). Our transformation is obtained from the transformation of Bassily and Smith with the following two modifications:

- (1) The transformation of Bassily and Smith applies to any “sample resilient” ϵ -LDP protocol. Informally, a sample resilient protocol is one whose outcome on any (distributed) database S is well-approximated by its outcome on a random subset of (roughly) half of the users in S . While we do not know of any protocols which are not sample resilient, we would like to obtain a more general result which avoids this restriction. We show that this restriction can be removed at the expense of increasing per-user communication to $O(\log \log n)$, where n is the number of users.
- (2) We generalize the transformation so that it also holds for any (ϵ, δ) -LDP protocol. Somewhat surprisingly, instead of obtaining an (ϵ, δ) -LDP protocol with short user reports, we show that, when done carefully, the resulting protocol actually satisfies pure $O(\epsilon)$ -privacy.

Our generic transformation is given in algorithm GenProt.

THEOREM 6.1. Let $\epsilon \leq \frac{1}{4}$, let $5 \ln(\frac{1}{\epsilon}) \leq T \leq \frac{1-e^{-\epsilon}}{4\delta e^\epsilon n}$, and let $\mathcal{M}$ be an (ϵ, δ) -LDP protocol. Then Algorithm GenProt satisfies 10ϵ -LDP. Moreover, for every database $S = (x_1, \dots, x_n)$, the total variation

distance between $\text{GenProt}(S)$ and $\mathcal{M}(S)$ is at most

$$\sup_F |\Pr[\text{GenProt}(S) \in F] - \Pr[\mathcal{M}(S) \in F]| \leq n \left(\left(\frac{1}{2} + \epsilon \right)^T + \frac{6T\delta e^\epsilon}{1 - e^{-\epsilon}} \right).$$

To interpret the guarantee of this theorem, let $\beta > 0$ be a parameter, and suppose $\delta \leq \epsilon\beta/48n \ln(2n/\beta)$. Then setting $T = 2 \ln(2n/\beta)$ makes this total variation distance at most β . With this setting of parameters, observe that each user must send $O(\log \log n)$ bits to the server. The analysis is omitted from this extended abstract.

7 A LOWER BOUND VIA ANTI-CONCENTRATION

In Section 3 we presented a LDP heavy-hitters algorithm with error

$$O\left(\frac{1}{\epsilon} \sqrt{n \cdot \log\left(\frac{|X|}{\beta}\right)}\right).$$

Ignoring the dependence on β , this error is known to be optimal. Specifically, Bassily and Smith [3] proved the following lower bound:

THEOREM 7.1 ([3]). Let $\epsilon = O(1)$ and $\delta = o\left(\frac{1}{n \log n}\right)$. Every non-interactive (ϵ, δ) -LDP protocol for estimating the frequencies of elements from a domain X must have worst-case error $\Omega\left(\frac{1}{\epsilon} \sqrt{n \cdot \log |X|}\right)$ with constant probability.

In this section we incorporate a tight dependence on the failure probability β into the lower bound. Specifically, we show

THEOREM 7.2. Let $\epsilon = O(1)$ and $\delta = o\left(\frac{1}{n \log n}\right)$. Every non-interactive (ϵ, δ) -LDP protocol for estimating the frequencies of elements from a domain X achieving worst-case error Δ with probability at least $1 - \beta$ must have

$$\Delta \geq \Omega\left(\frac{1}{\epsilon} \sqrt{n \cdot \log\left(\frac{|X|}{\beta}\right)}\right).$$

In light of Theorem 7.1, we can fix $X = \{0, 1\}$ and only show a lower bound of $\Omega\left(\frac{1}{\epsilon} \sqrt{n \cdot \log(\frac{1}{\beta})}\right)$. To that end, we strengthen an argument of Chan et al. [5] (see also [14, 26]) who obtained a lower bound of $\Omega(\sqrt{n})$ on the error:

THEOREM 7.3 ([5, 14, 26]). Let $\epsilon \leq 1/2$ and $\delta < 1/4n$. Every (ϵ, δ) -LDP frequency oracle must have worst-case error $\Omega(\sqrt{n})$ with constant probability.

We now present a high-level overview of the proof of Theorem 7.2. In this overview, we argue how the result follows from the ideas presented in the previous sections; namely, from advanced

Algorithm 3 GenProt

Inputs: User's inputs $\{x_i \in X : i \in [n]\}$, privacy parameter ϵ , and parameter $T \in \mathbb{N}$.

Algorithms used: An (ϵ, δ) -LDP protocol $\mathcal{M}$ with local randomizers $\mathcal{A}_i, i \in [n]$.

1. For every $(i, t) \in [n] \times [T]$ generate independent public string: $y_{i,t} \leftarrow \mathcal{A}_i(\perp)$.

2. For user $i = 1$ to n do

(a) For every $t \in [T]$ compute $p_{i,t} = \frac{1}{2} \frac{\Pr[\mathcal{A}_i(x_i)=y_{i,t}]}{\Pr[\mathcal{A}_i(\perp)=y_{i,t}]}$.

(b) For every $t \in [T]$, if $p_{i,t} \notin \left[\frac{e^{-2\epsilon}}{2}, \frac{e^{2\epsilon}}{2}\right]$ then set $p_{i,t} = \frac{1}{2}$.

(c) For every $t \in [T]$ sample a bit $b_{i,t}$ from Bernoulli($p_{i,t}$).

(d) Denote $H_i = \{t \in [T] : b_{i,t} = 1\}$.

(e) If $H_i = \emptyset$ then set $H_i = [T]$.

(f) Sample $g_i \in H_i$ uniformly, and send g_i to the server.

3. Run $\mathcal{M}(y_{1,g_1}, \dots, y_{n,g_n})$.

groupposition and from our generic transformation from approximate to pure LDP. However, in the full version of this work we present the proof in a somewhat more streamlined way that does not go through these generic results.

By our transformation from approximate to pure LDP, we may assume without loss of generality that we have a pure ϵ -LDP protocol $\mathcal{A}$ for n users, which for every (distributed) database $D \in \{0, 1\}^n$ counts the number of 1's in D to within error Δ with success probability $1 - \beta$. Our goal is to show that $\Delta = \Omega\left(\frac{1}{\epsilon} \sqrt{n \cdot \log\left(\frac{1}{\beta}\right)}\right)$.

We generate an input to $\mathcal{A}$ as follows: For some constant $C > 0$, denote $m = C\epsilon^2 n$, and let $S = (X_1, \dots, X_m) \in \{0, 1\}^m$ be a database chosen uniformly at random. Now define $D = (Y_1, \dots, Y_n) \in \{0, 1\}^n$, where $Y_i = X_{\lceil im/n \rceil}$. That is, the first $\frac{m}{n}$ elements in D are X_1 , the next $\frac{m}{n}$ elements are X_2 , and so on. Now consider applying the protocol $\mathcal{A}$ onto the database D . The remainder of the proof proceeds as follows:

- (i) On one hand, the error of $\mathcal{A}$ on D is at most Δ , and hence, renormalizing to S , we obtain an estimate of the number of 1's in S with error at most $\frac{m}{n} \Delta = C\epsilon^2 \Delta$.
- (ii) On the other hand, every bit in S has $\frac{n}{m} = \frac{1}{C\epsilon^2}$ copies in D . By advanced groupposition, we provide every element in S with say $(\frac{1}{100}, \frac{1}{100})$ -DP. Using bounds on the mutual information of a random input and the output of a DP algorithm [3], we get that (a lot of) the elements in S remain approximately uniform even conditioned on the transcript of the protocol (and they also remain independent).

Combining (i) and (ii), and applying an anti-concentration bound to the sum of elements in S , we will get that

$$C\epsilon^2 \Delta \gtrsim \sqrt{|S| \cdot \log\left(\frac{1}{\beta}\right)} = \sqrt{C\epsilon^2 n \cdot \log\left(\frac{1}{\beta}\right)},$$

which gives the result. The full proof of Theorem 7.2 is omitted from this extended abstract.

Acknowledgements. We thank Kobbi Nissim for suggesting the name “advanced groupposition” for the phenomenon described in Section 4.

REFERENCES

- [1] Noga Alon and Joel Spencer. 1992. *The Probabilistic Method*. John Wiley.
- [2] Raef Bassily, Kobbi Nissim, Uri Stemmer, and Abhradeep Thakurta. 2017. Practical Locally Private Heavy Hitters. In *Advances in Neural Information Processing Systems (NIPS)*.
- [3] Raef Bassily and Adam D. Smith. 2015. Local, Private, Efficient Protocols for Succinct Histograms. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*. 127–135. <https://doi.org/10.1145/2746539.2746632>
- [4] Mark Bun and Thomas Steinke. 2016. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*. 635–658.
- [5] T.-H. Hubert Chan, Elaine Shi, and Dawn Song. 2012. Optimal Lower Bound for Differentially Private Multi-party Aggregation. In *Algorithms - ESA 2012 - 20th Annual European Symposium, Ljubljana, Slovenia, September 10-12, 2012. Proceedings*. 277–288. https://doi.org/10.1007/978-3-642-33090-2_25
- [6] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. 2013. Local Privacy and Statistical Minimax Rates. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*. 429–438. <https://doi.org/10.1109/FOCS.2013.53>
- [7] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. 2015. Generalization in Adaptive Data Analysis and Holdout Reuse. In *Advances in Neural Information Processing Systems (NIPS)*. Montreal.
- [8] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC (Lecture Notes in Computer Science)*, Vol. 3876. Springer, 265–284.
- [9] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3-4 (2014), 211–407. <https://doi.org/10.1561/04000000042>
- [10] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. 2010. Boosting and Differential Privacy. In *FOCS. IEEE Computer Society*, 51–60.
- [11] Ulfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *CCS*.
- [12] Anna C. Gilbert, Yi Li, Ely Porat, and Martin J. Strauss. 2014. For-All Sparse Recovery in Near-Optimal Time. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP)*. 538–550.
- [13] V. Guruswami and P. Indyk. 2001. Expander-based constructions of efficiently decodable codes. In *Proceedings 2001 IEEE International Conference on Cluster Computing*. 658–667. <https://doi.org/10.1109/SFCS.2001.959942>
- [14] Justin Hsu, Sanjeev Khanna, and Aaron Roth. 2012. Distributed Private Heavy Hitters. In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I*. 461–472. https://doi.org/10.1007/978-3-642-31594-7_39
- [15] Daniel M. Kane, Jelani Nelson, Ely Porat, and David P. Woodruff. 2011. Fast moment estimation in data streams in optimal space. In *STOC. ACM*, 745–754.
- [16] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2011. What Can We Learn Privately? *SIAM J. Comput.* 40, 3 (2011), 793–826.
- [17] K. G. Larsen, J. Nelson, H. L. Nguyen, and M. Thorup. 2016. Heavy Hitters via Cluster-Preserving Clustering. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. 61–70. <https://doi.org/10.1109/FOCS.2016.16>

- [18] Nina Mishra and Mark Sandler. 2006. Privacy via Pseudorandom Sketches. In *Proceedings of the Twenty-fifth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS '06)*. ACM, New York, NY, USA, 143–152. <https://doi.org/10.1145/1142351.1142373>
- [19] Michael Mitzenmacher and Eli Upfal. 2005. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, New York, NY, USA.
- [20] Kobbi Nissim and Uri Stemmer. 2017. Clustering Algorithms for the Centralized and Local Models. (2017). [arXiv.org](https://arxiv.org/).
- [21] Zhan Qin, Yin Yang, Ting Yu, Issa Khalil, Xiaokui Xiao, and Kui Ren. 2016. Heavy Hitter Estimation over Set-Valued Data with Local Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*. ACM, New York, NY, USA, 192–203. <https://doi.org/10.1145/2976749.2978409>
- [22] Ryan M. Rogers, Aaron Roth, Adam D. Smith, and Om Thakkar. 2016. Max-Information, Differential Privacy, and Post-selection Hypothesis Testing. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*. 487–494.
- [23] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. 1995. Chernoff-Hoeffding Bounds for Applications with Limited Independence. *SIAM Journal on Discrete Mathematics* 8, 2 (1995), 223–250. <https://doi.org/10.1137/S089548019223872X> arXiv:<https://doi.org/10.1137/S089548019223872X>
- [24] A. Smith, A. Thakurta, and J. Upadhyay. 2017. Is Interaction Necessary for Distributed Private Learning?. In *2017 IEEE Symposium on Security and Privacy (SP)*. 58–77. <https://doi.org/10.1109/SP.2017.35>
- [25] A.G. Thakurta, A.H. Vyrros, U.S. Vaishampayan, G. Kapoor, J. Freudiger, V.R. Sridhar, and D. Davidson. 2017. Learning new words. *US Patent 9594741* (2017).
- [26] Salil Vadhan. 2016. *The Complexity of Differential Privacy*.