Between Pure and Approximate Differential Privacy

Thomas Steinke* Jonathan Ullman†

January 24, 2015

Abstract

We show a new lower bound on the sample complexity of (ε, δ)-differentially private algorithms that accurately answer statistical queries on high-dimensional databases. The novelty of our bound is that it depends optimally on the parameter δ, which loosely corresponds to the probability that the algorithm fails to be private, and is the first to smoothly interpolate between approximate differential privacy ($\delta > 0$) and pure differential privacy ($\delta = 0$).

Specifically, we consider a database $D \in \{\pm 1\}^{n \times d}$ and its *one-way marginals*, which are the d queries of the form "What fraction of individual records have the i-th bit set to +1?" We show that in order to answer all of these queries to within error $\pm\alpha$ (on average) while satisfying (ε, δ)-differential privacy, it is necessary that

$$n \geq \Omega\left(\frac{\sqrt{d \log(1/\delta)}}{\alpha\varepsilon}\right),$$

which is optimal up to constant factors. To prove our lower bound, we build on the connection between *fingerprinting codes* and lower bounds in differential privacy (Bun, Ullman, and Vadhan, STOC'14).

In addition to our lower bound, we give new purely and approximately differentially private algorithms for answering arbitrary statistical queries that improve on the sample complexity of the standard Laplace and Gaussian mechanisms for achieving worst-case accuracy guarantees by a logarithmic factor.

*Harvard University School of Engineering and Applied Sciences. Supported by NSF grant CCF-1116616. Email: tsteinke@seas.harvard.edu.

†Columbia University Department of Computer Science. Supported by a Junior Fellowship from the Simons Society of Fellows. Email: jullman@cs.columbia.edu.

Contents

1 Introduction

The goal of privacy-preserving data analysis is to enable rich statistical analysis of a database while protecting the privacy of individuals whose data is in the database. A formal privacy guarantee is given by (ε, δ)-*differential privacy* [DMNS06, DKM+06], which ensures that no individual's data has a significant influence on the information released about the database. The two parameters ε and δ control the level of privacy. Very roughly, ε is an upper bound on the amount of influence an individual's record has on the information released and δ is the probability that this bound fails to hold[1], so the definition becomes more stringent as $\varepsilon, \delta \to 0$.

A natural way to measure the tradeoff between privacy and utility is *sample complexity*—the minimum number of records n that is sufficient in order to publicly release a given set of statistics about the database, while achieving both differential privacy and statistical accuracy. Intuitively, it's easier to achieve these two goals when n is large, as each individual's data will have only a small influence on the aggregate statistics of interest. Conversely, the sample complexity n should increase as ε and δ decrease (which strengthens the privacy guarantee).

The strongest version of differential privacy, in which $\delta = 0$, is known as *pure differential privacy*. The sample complexity of achieving pure differential privacy is well known for many settings (e.g. [HT10]). The more general case where $\delta > 0$ is known as *approximate differential privacy*, and is less well understood. Recently, Bun, Ullman, and Vadhan [BUV14] showed how to prove strong lower bounds for approximate differential privacy that are essentially optimal for $\delta \approx 1/n$, which is essentially the weakest privacy guarantee that is still meaningful.[2]

Since δ bounds the probability of a complete privacy breach, we would like δ to be very small. Thus we would like to quantify the cost (in terms of sample complexity) as $\delta \to 0$. In this work we give lower bounds for approximately differentially private algorithms that are nearly optimal for every choice of δ, and smoothly interpolate between pure and approximate differential privacy.

Specifically, we consider algorithms that compute the *one-way marginals of the database*—an extremely simple and fundamental family of queries. For a database $D \in \{\pm 1\}^{n \times d}$, the d one-way marginals are simply the mean of the bits in each of the d columns. Formally, we define

$$\overline{D} := \frac{1}{n} \sum_{i=1}^{n} D_i \in [\pm 1]^d$$

where $D_i \in \{\pm 1\}^d$ is the i-th row of D. A mechanism M is said to be *accurate* if, on input D, its output is "close to" $\overline{D}$. Accuracy may be measured in a *worst-case* sense—i.e. $\left\| M(D) - \overline{D} \right\|_\infty \leq \alpha$, meaning every one-way marginal is answered with accuracy α—or in an *average-case* sense—i.e. $\left\| M(D) - \overline{D} \right\|_1 \leq \alpha d$, meaning the marginals are answered with average accuracy α.

Some of the earliest results in differential privacy [DN03, DN04, BDMN05, DMNS06] give a simple (ε, δ)-differentially private algorithm—the *Laplace mechanism*—that computes the one-way marginals of $D \in \{\pm 1\}^{n \times d}$ with average error α as long as

$$n \geq O\left(\min\left\{ \frac{\sqrt{d \log(1/\delta)}}{\varepsilon \alpha}, \frac{d}{\varepsilon \alpha} \right\} \right). \tag{1}$$

[1] This intuition is actually somewhat imprecise, although it is suitable for this informal discussion. See [KS08] for a more precise semantic interpretation of (ε, δ)-differential privacy.

[2] When $\delta \geq 1/n$ there are algorithms that are intuitively not private, yet satisfy $(0, \delta)$-differential privacy.

The previous best lower bounds are $n \geq \Omega(d/\varepsilon\alpha)$ [HT10] for pure differential privacy and $n \geq \tilde{\Omega}(\sqrt{d}/\varepsilon\alpha)$ for approximate differential privacy with $\delta = o(1/n)$ [BUV14]. Our main result is an optimal lower bound that combines the previous lower bounds.

Theorem 1.1 (Main Theorem)**.** *For every $\varepsilon \leq O(1)$, every $2^{-\Omega(n)} \leq \delta \leq 1/n^{1+\Omega(1)}$ and every $\alpha \leq 1/10$, if $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$ is (ε, δ)-differentially private and $\mathop{\mathbb{E}}_{M}\left[\|M(D) - \overline{D}\|_1\right] \leq \alpha d$, then*

$$n \geq \Omega\left(\frac{\sqrt{d \log(1/\delta)}}{\varepsilon\alpha}\right).$$

More generally, this is the first result showing that the sample complexity must grow by a multiplicative factor of $\sqrt{\log(1/\delta)}$ for answering any family of queries, as opposed to an additive dependence on δ. We also remark that the assumption on the range of δ is necessary, as the Laplace mechanism gives accuracy α and satisfies $(\varepsilon, 0)$-differential privacy when $n \geq O(d/\varepsilon\alpha)$.

1.1 Average-Case Versus Worst-Case Error

Our lower bound holds for mechanisms with an average-case (L_1) error guarantee. Thus, it also holds for algorithms that achieve worst-case (L_∞) error guarantees. The Laplace mechanism gives a matching upper bound for average-case error. In many cases worst-case error guarantees are preferrable. For worst-case error, the sample complexity of the Laplace mechanism degrades by an additional $\log d$ factor compared to (1).

Surprisingly, this degradation is not necessary. We present algorithms that answer every one-way marginal with α accuracy and improve on the sample complexity of the Laplace mechanism by roughly a $\log d$ factor. These algorithms demonstrate that the widely used technique of adding independent noise to each query is suboptimal when the goal is to achieve worst-case error guarantees.

Our algorithm for pure differential privacy satisfies the following.

Theorem 1.2. *For every $\varepsilon, \alpha > 0$, $d \geq 1$, and $n \geq 4d/\varepsilon\alpha$, there exists an efficient mechanism $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$ that is $(\varepsilon, 0)$-differentially private and*

$$\forall D \in \{\pm 1\}^{n \times d} \quad \mathop{\mathbb{P}}_{M}\left[\left\|M(D) - \overline{D}\right\|_\infty \geq \alpha\right] \leq (2e)^{-d}.$$

And our algorithm for approximate differential privacy is as follows.

Theorem 1.3. *For every $\varepsilon, \delta, \alpha > 0$, $d \geq 1$, and*

$$n \geq O\left(\frac{\sqrt{d \cdot \log(1/\delta) \cdot \log\log d}}{\varepsilon\alpha}\right),$$

there exists an efficient mechanism $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$ that is (ε, δ)-differentially private and

$$\forall D \in \{\pm 1\}^{n \times d} \quad \mathop{\mathbb{P}}_{M}\left[\left\|M(D) - \overline{D}\right\|_\infty \geq \alpha\right] \leq \frac{1}{d^{\omega(1)}}.$$

These algorithms improve over the sample complexity of the best known mechanisms for each privacy and accuracy guarantee by a factor of $(\log(d))^{\Omega(1)}$. Namely, the Laplace mechanism requires $n \geq O(d \cdot \log d/\varepsilon\alpha)$ samples for pure differential privacy and the Gaussian mechanism requires $n \geq O(\sqrt{d \cdot \log(1/\delta) \cdot \log d}/\varepsilon\alpha)$ samples for approximate differential privacy.

Privacy	Accuracy	Type	Previous bound	This work
(ε, δ)	L_1 or L_∞	Lower	$n = \tilde{\Omega}\left(\frac{\sqrt{d}}{\alpha\varepsilon}\right)$ [BUV14]	$n = \Omega\left(\frac{\sqrt{d\log(1/\delta)}}{\alpha\varepsilon}\right)$
(ε, δ)	L_1	Upper	$n = O\left(\frac{\sqrt{d\cdot\log(1/\delta)}}{\varepsilon\alpha}\right)$ Laplace	
(ε, δ)	L_∞	Upper	$n = O\left(\frac{\sqrt{d\cdot\log(1/\delta)\cdot\log d}}{\varepsilon\alpha}\right)$ Gaussian	$n = O\left(\frac{\sqrt{d\cdot\log(1/\delta)\cdot\log\log d}}{\varepsilon\alpha}\right)$
ε	L_1 or L_∞	Lower	$n = \Omega\left(\frac{d}{\alpha\varepsilon}\right)$ [HT10]	
ε	L_1	Upper	$n = O\left(\frac{d}{\varepsilon\alpha}\right)$ Laplace	
ε	L_∞	Upper	$n = O\left(\frac{d\cdot\log d}{\varepsilon\alpha}\right)$ Laplace	$n = O\left(\frac{d}{\varepsilon\alpha}\right)$

Figure 1: Summary of sample complexity upper and lower bounds for privately answering d one-way marginals with accuracy α.

1.2 Techniques

Lower Bounds: Our lower bound builds on the work of Bun, Ullman, and Vadhan [BUV14]. Their work uses a combinatorial object called a *fingerprinting code* [BS98] in a black-box manner to prove a lower bound for one-way marginals. Fingerprinting codes were originally used in cryptography for watermarking digital content; however, they showed that they can be used to construct an attack demonstrating that any mechanism that accurately answers one-way marginals is not differentially private. In particular, a fingerprinting code gives a distribution on individuals' data and a corresponding "tracer" algorithm such that, if a database is constructed from the data of a fixed subset of the individuals, then the tracer algorithm can identify at least one of the individuals in that subset given only approximate answers to the one-way marginals of the database. Specifically, their attack shows that a mechanism that satisfies $(1, o(1/n))$-differential privacy requires $n \geq \tilde{\Omega}(\sqrt{d})$ samples to accurately compute one-way marginals.

Our proof uses a new, more general reduction from breaking fingerprinting codes to differentially private data release. Specifically, our reduction uses *group differential privacy*. This property states that if an algorithm is (ε, δ)-differentially private with respect to the change of one individual's data, then for any k, it is roughly $(k\varepsilon, e^{k\varepsilon}\delta)$-differentially private with respect to the change of k individuals' data. Thus an (ε, δ)-differentially private algorithm provides a meaningful privacy guarantee for groups of size $k \approx \log(1/\delta)/\varepsilon$.

To use this in our reduction, we start with a mechanism M that takes a database of n rows and is (ε, δ)-differentially private. We design a mechanism M_k that takes a database of n/k rows, copies each of its rows k times, and uses the result as input to M. The resulting mechanism M_k is roughly $(k\varepsilon, e^{k\varepsilon}\delta)$-differentially private. For our choice of k, these parameters will be small enough to apply the attack of [BUV14] to obtain a lower bound on the number of samples used by M_k, which is n/k. Thus, for larger values of k (equivalently, smaller values of δ), we obtain a stronger lower bound. The remainder of the proof is to quantify the parameters precisely.

Upper Bounds: Our algorithm for pure differential privacy and worst-case error is an instantiation of the exponential mechanism [MT07] using the L_∞ norm. That is, the mechanism samples $y \in \mathbb{R}^d$ with probability proportional to $\exp(-\eta \left\|y\right\|_\infty)$ and outputs $M(D) = \overline{D} + y$. In

contrast, adding independent Laplace noise corresponds to using the exponential mechanism with the L_1 norm and adding independent Gaussian noise corresponds to using the exponential mechanism with the L_2 norm squared. Using this distribution turns out to give better tail bounds than adding independent noise.

For approximate differential privacy, we use a completely different algorithm. We start by adding independent Gaussian noise to each marginal. However, rather than using a union bound to show that each Gaussian error is small with high probability, we use a Chernoff bound to show that most errors are small. Namely, with the sample complexity that we allow M, we can ensure that all but a $1/\text{polylog}(d)$ fraction of the errors are small. Now we "fix" the $d/\text{polylog}(d)$ marginals that are bad. The trick is that we use the sparse vector algorithm, which allows us to do indentify and fix these $d/\text{polylog}(d)$ marginals with sample complexity corresponding to only $d/\text{polylog}(d)$ queries, rather than d queries.

2 Preliminaries

We define a *database* $D \in \{\pm 1\}^{n \times d}$ to be a matrix of n rows, where each row corresponds to an individual, and each row has *dimension* d (consists of d binary attributes). We say that two databases $D, D' \in \{\pm 1\}^{n \times d}$ are *adjacent* if they differ only by a single row, and we denote this by $D \sim D'$. In particular, we can replace the ith row of a database D with some fixed element of $\{\pm 1\}^d$ to obtain another database $D_{-i} \sim D$.

Definition 2.1 (Differential Privacy [DMNS06])**.** Let $M : \{\pm 1\}^{n \times d} \to \mathcal{R}$ be a randomized mechanism. We say that M is (ε, δ)-*differentially private* if for every two adjacent databases $D \sim D'$ and every subset $S \subseteq \mathcal{R}$,

$$\mathbb{P}[M(D) \in S] \leq e^{\varepsilon} \cdot \mathbb{P}\left[M(D') \in S\right] + \delta.$$

A well known fact about differential privacy is that it generalizes smoothly to databases that differ on more than a single row. We say that two databases $D, D' \in \{\pm 1\}^{n \times d}$ are k-*adjacent* if they differ by at most k rows, and we denote this by $D \sim_k D'$.

Fact 2.2 (Group Differential Privacy)**.** *For every $k \geq 1$, if $M : \{\pm 1\}^{n \times d} \to \mathcal{R}$ is (ε, δ)-differentially private, then for every two k-adjacent databases $D \sim_k D'$, and every subset $S \subseteq \mathcal{R}$,*

$$\mathbb{P}[M(D) \in S] \leq e^{k\varepsilon} \cdot \mathbb{P}\left[M(D') \in S\right] + \frac{e^{k\varepsilon} - 1}{e^{\varepsilon} - 1} \cdot \delta.$$

All of the upper and lower bounds for one-way marginals have a multiplicative $1/\alpha\varepsilon$ dependence on the accuracy α and the privacy loss ε. This is no coincidence - there is a generic reduction:

Fact 2.3 (α and ε dependence)**.** *Let $p \in [1, \infty]$ and $\alpha, \varepsilon, \delta \in [0, 1/10]$.*

Suppose there exists a (ε, δ)-differentially private mechanism $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$ such that for every database $D \in \{\pm 1\}^{n \times d}$,

$$\mathop{\mathbb{E}}_{M}\left[\left\|M(D) - \overline{D}\right\|_p\right] \leq \alpha d^{1/p}.$$

Then there exists a $(1, \delta/\varepsilon)$-differentially private mechanism $M' : \{\pm 1\}^{n' \times d} \to [\pm 1]^d$ for $n' = \Theta(\alpha\varepsilon n)$ such that for every database $D' \in \{\pm 1\}^{n' \times d}$,

$$\mathop{\mathbb{E}}_{M'}\left[\left\|M'(D') - \overline{D'}\right\|_p\right] \leq d^{1/p}/10.$$

This fact allows us to suppress the accuracy parameter α and the privacy loss ε when proving our lower bounds. Namely, if we prove a lower bound of $n' \geq n^*$ for all $(1,\delta)$-differentially private mechanisms $M' : \{\pm 1\}^{n' \times d} \to [\pm 1]^d$ with $\mathop{\mathbb{E}}_{M'}\left[\|M'(D') - \overline{D'}\|_p\right] \leq d^{1/p}/10$, then we obtain a lower bound of $n \geq \Omega(n^*/\alpha\varepsilon)$ for all $(\varepsilon, \varepsilon\delta)$-differentially private mechanisms $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$ with $\mathop{\mathbb{E}}_{M}\left[\|M(D) - \overline{D}\|_p\right] \leq \alpha d^{1/p}$. So we will simply fix the parameters $\alpha = 1/10$ and $\varepsilon = 1$ in our lower bounds.

3 Lower Bounds for Approximate Differential Privacy

Our main theorem can be stated as follows.

Theorem 3.1 (Main Theorem)**.** *Let $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$ be a $(1,\delta)$-differentially private mechanism that answers one-way marginals such that*

$$\forall D \in \{\pm 1\}^{n \times d} \quad \mathop{\mathbb{E}}_{M}\left[\left\|M(D) - \overline{D}\right\|_1\right] \leq \frac{d}{10},$$

where $\overline{D}$ is the true answer vector. If $2^{-Omega(n)} \leq \delta \leq 1/n^{1+\Omega(1)}$ and n is sufficiently large, then

$$d \leq O\left(\frac{n^2}{\log(1/\delta)}\right).$$

Theorem 1.1 in the introduction follows by rearranging terms, and applying Fact 2.3. The statement above is more convenient technically, but the statement in the introduction is more consistent with the literature.

First we must introduce fingerprinting codes. The following definition is tailored to the application to privacy. Fingerprinting codes were originally defined by Boneh and Shaw [BS98] with a worst-case accuracy guarantee. Subsequent works [BUV14, SU14] have altered the accuracy guarantee to an average-case one, which we use here.

Definition 3.2 (L_1 Fingerprinting Code)**.** A *ε-complete δ-sound α-robust L_1 fingerprinting code for n users with length d* is a pair of random variables $D \in \{\pm 1\}^{n \times d}$ and $Trace : [\pm 1]^d \to 2^{[n]}$ such that the following hold.

Completeness: For any fixed $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$,

$$\mathbb{P}\left[\left(\left\|M(D) - \overline{D}\right\|_1 \leq \alpha d\right) \wedge (Trace(M(D)) = \emptyset)\right] \leq \varepsilon.$$

Soundness: For any $i \in [n]$ and fixed $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$,

$$\mathbb{P}\left[i \in Trace(M(D_{-i}))\right] \leq \delta,$$

where D_{-i} denotes D with the i^{th} row replaced by some fixed element of $\{\pm 1\}^d$.

Fingerprinting codes with optimal length were first constructed by Tardos [Tar08] (for worst-case error) and subsequent works [BUV14, SU14] have adapted Tardos' construction to work for average-case error guarantees, which yields the following theorem.

Theorem 3.3. *For every $n \geq 1$, $\delta > 0$, and $d \geq d_{n,\delta} = O(n^2 \log(1/\delta))$, there exists a 1/100-complete δ-sound 1/8-robust L_1 fingerprinting code for n users with length d.*

We now show how the existence of fingerprinting codes implies our lower bound.

Proof of Theorem 3.1 from Theorem 3.3. Let $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$ be a $(1, \delta)$-differentially private mechanism such that

$$\forall D \in \{\pm 1\}^{n \times d} \quad \mathop{\mathbb{E}}_{M}\left[\left\|M(D) - \overline{D}\right\|_1\right] \leq \frac{d}{10}.$$

Then, by Markov's inequality,

$$\forall D \in \{\pm 1\}^{n \times d} \quad \mathop{\mathbb{P}}_{M}\left[\left\|M(D) - \overline{D}\right\|_1 > \frac{d}{9}\right] \leq \frac{9}{10}. \tag{2}$$

Let k be a parameter to be chosen later. Let $n_k = \lfloor n/k \rfloor$. Let $M_k : \{\pm 1\}^{n_k \times d} \to [\pm 1]^d$ be the following mechanism. On input $D^* \in \{\pm 1\}^{n_k \times d}$, M_k creates $D \in \{\pm 1\}^{n \times d}$ by taking k copies of D^* and filling the remaining entries with 1s. Then M_k runs M on D and outputs $M(D)$.

By group privacy (Fact 2.2), M_k is a $\left(\varepsilon_k = k, \delta_k = \frac{e^k - 1}{e - 1}\delta\right)$-differentially private mechanism. By the triangle inequality,

$$\left\|M_k(D^*) - \overline{D^*}\right\|_1 \leq \left\|M(D) - \overline{D}\right\|_1 + \left\|\overline{D} - \overline{D^*}\right\|_1. \tag{3}$$

Now

$$\overline{D}_j = \frac{k \cdot n_k}{n}\overline{D^*_j} + \frac{n - k \cdot n_k}{n} 1.$$

Thus

$$\left|\overline{D}_j - \overline{D^*_j}\right| = \left|\left(\frac{k \cdot n_k}{n} - 1\right)\overline{D^*_j} + \frac{n - k \cdot n_k}{n}\right| = \frac{n - k \cdot n_k}{n}\left|1 - \overline{D^*_j}\right| \leq 2\frac{n - k \cdot n_k}{n}.$$

We have

$$\frac{n - k \cdot n_k}{n} = \frac{n - k\lfloor n/k \rfloor}{n} \leq \frac{n - k(n/k - 1)}{n} = \frac{k}{n}.$$

Thus $\left\|\overline{D} - \overline{D^*}\right\|_1 \leq 2k/n$. Assume $k \leq n/100$. Thus $\left\|\overline{D} - \overline{D^*}\right\|_1 \leq d/50$ and, by (2) and (3),

$$\mathop{\mathbb{P}}_{M_k}\left[\left\|M_k(D^*) - \overline{D^*}\right\|_1 > \frac{d}{8}\right] \leq \mathop{\mathbb{P}}_{M}\left[\left\|M(D) - \overline{D}\right\|_1 > \frac{d}{9}\right] \leq \frac{9}{10}. \tag{4}$$

Assume $d \geq d_{n_k,\delta}$, were $d_{n_k,\delta} = O(n_k^2 \log(1/\delta))$ is as in Theorem 3.3. We will show by contradiction that this cannot be – that is $d \leq O(n_k^2 \log(1/\delta))$. Let $D^* \in \{\pm 1\}^{n_k \times d}$ and $Trace : [\pm 1]^d \to 2^{[n_k]}$ be a 1/100-complete δ-sound 1/8-robust L_1 fingerprinting code for n_k users of length d.

By the completeness of the fingerprinting code,

$$\mathbb{P}\left[\left\|M_k(D^*) - \overline{D^*}\right\|_1 \leq \frac{d}{8} \wedge Trace(M(D)) = \emptyset\right] \leq \frac{1}{100}. \tag{5}$$

Combinging (4) and (5), gives

$$\mathbb{P}\left[Trace(M_k(D^*)) \neq \emptyset\right] \geq \frac{9}{100} > \frac{1}{12}.$$

In particular, there exists $i^* \in [n_k]$ such that

$$\mathbb{P}[i^* \in Trace(M_k(D^*))] \geq > \frac{1}{12n_k}. \tag{6}$$

We have that $Trace(M_k(D^*))$ is a $(\varepsilon_k, \delta_k)$-differentially private function of D^*, as it is only postprocessing $M_k(D^*)$. Thus

$$\mathbb{P}[i^* \in Trace(M_k(D^*))] \leq e^{\varepsilon_k}\mathbb{P}\left[i^* \in Trace(M_k(D^*_{-i^*}))\right] + \delta_k \leq e^{\varepsilon_k}\delta + \delta_k, \tag{7}$$

where the second inequality follows from the soundness of the fingerprinting code.

Combining (6) and (7) gives

$$\frac{1}{12n_k} \leq e^{\varepsilon_k}\delta + \delta_k = e^k\delta + \frac{e^k - 1}{e-1}\delta = \frac{e^{k+1}-1}{e-1}\delta < e^{k+1}\delta. \tag{8}$$

If $k \leq \log(1/12n_k\delta) - 1$, then (8) gives a contradiction. Let $k = \lfloor \log(1/12n\delta) - 1 \rfloor$. Assuming $\delta \geq e^{-n/100}$ ensures $k \leq n/100$, as required. Assuming $\delta \leq 1/n^{1+\gamma}$ implies $k \geq \log(1/\delta)/(1+1/\gamma) - 5 \geq \Omega(\log(1/\delta))$. This setting of k gives a contradiction, which implies that

$$d < d_{n_k,\delta} = O(n_k^2 \log(1/\delta)) = O\left(\frac{n^2}{k^2}\log(1/\delta)\right) = O\left(\frac{n^2}{\log(1/\delta)}\right),$$

as required.

□

4 New Mechanisms for L_∞ Error

Adding independent noise seems very natural for one-way marginals, but it is suboptimal if one is interested in worst-case (i.e. L_∞) error bounds, rather than average-case (i.e. L_1) error bounds.

4.1 Pure Differential Privacy

Theorem 1.2 follows from Theorem 4.1. In particular, the mechanism $M : \{\pm 1\}^{n \times d} \to [\pm 1]^d$ in Theorem 1.2 is given by $M(D) = \overline{D} + Y$, where $Y \sim \mathcal{D}$ and $\mathcal{D}$ is the distribution from Theorem 4.1 with $\Delta = 2/n$.[3]

Theorem 4.1. *For all $\varepsilon > 0$, $d \geq 1$, and $\Delta > 0$, there exists a continuous distribution $\mathcal{D}$ on $\mathbb{R}^d$ with the following properties.*

- ***Privacy:*** *If $x, x' \in \mathbb{R}^d$ with $\|x - x'\|_\infty \leq \Delta$, then*

$$\mathbb{P}_{Y\sim\mathcal{D}}[x + Y \in S] \leq e^\varepsilon \mathbb{P}_{Y\sim\mathcal{D}}[x' + Y \in S]$$

for all measurable $S \subseteq \mathbb{R}^d$.

[3]Note that we must truncate the output of M to ensure that $M(D)$ is always in $[\pm 1]^d$.

- ***Accuracy:*** *For all* $\alpha > 0$,

$$\mathop{\mathbb{P}}_{Y\sim\mathcal{D}}[\|Y\|_\infty \geq \alpha] \leq \left(\frac{\Delta d}{\varepsilon\alpha}\right)^d e^{d-\alpha\varepsilon/\Delta}.$$

  In particular, if $d \leq \varepsilon\alpha/2\Delta$, *then* $\mathop{\mathbb{P}}_{Y\sim\mathcal{D}}[\|Y\|_\infty \geq \alpha] \leq (2e)^{-d}$.

- ***Efficiency:*** $\mathcal{D}$ *can be efficiently sampled.*

Proof. The distribution $\mathcal{D}$ is simply an instantiation of the exponential mechanism [MT07]. In particular, the probability density function is given by

$$\mathrm{pdf}_{\mathcal{D}}(y) \propto \exp\left(-\frac{\varepsilon}{\Delta}\left\|y\right\|_\infty\right).$$

Formally, for every measurable $S \subseteq \mathbb{R}^d$,

$$\mathop{\mathbb{P}}_{Y\sim\mathcal{D}}[Y \in S] = \frac{\int_S \exp\left(-\frac{\varepsilon}{\Delta}\left\|y\right\|_\infty\right)\mathrm{d}y}{\int_{\mathbb{R}^d} \exp\left(-\frac{\varepsilon}{\Delta}\left\|y\right\|_\infty\right)\mathrm{d}y}.$$

Firstly, this is clearly a well-defined distribution as long as $\varepsilon/\Delta > 0$.

Privacy is easy to verify: It suffices to bound the ratio of the probability densities for the shifted distributions. For $x, x' \in \mathbb{R}^d$ with $\|x' - x\|_\infty \leq \Delta$, by the triangle inequality,

$$\frac{\mathrm{pdf}_{\mathcal{D}}(x+y)}{\mathrm{pdf}_{\mathcal{D}}(x'+y)} = \frac{\exp\left(-\frac{\varepsilon}{\Delta}\left\|x+y\right\|_\infty\right)}{\exp\left(-\frac{\varepsilon}{\Delta}\left\|x'+y\right\|_\infty\right)} = \exp\left(\frac{\varepsilon}{\Delta}\left(\left\|x'+y\right\|_\infty - \left\|x+y\right\|_\infty\right)\right) \leq \exp\left(\frac{\varepsilon}{\Delta}\left\|x'-x\right\|_\infty\right) \leq e^{\varepsilon}.$$

Define a distribution $\mathcal{D}^*$ on $[0,\infty)$ to by $Z \sim \mathcal{D}^*$ meaning $Z = \|Y\|_\infty$ for $Y \sim \mathcal{D}$. To prove accuracy, we must give a tail bound on $\mathcal{D}^*$. The probability density function of $\mathcal{D}^*$ is given by

$$\mathrm{pdf}_{\mathcal{D}^*}(z) \propto z^{d-1} \cdot \exp\left(-\frac{\varepsilon}{\Delta}z\right),$$

which is obtained by integrating the probability density function of $\mathcal{D}$ over the infinity-ball of radius z, which has surface area $d2^d z^{d-1} \propto z^{d-1}$. Thus $\mathcal{D}^*$ is precisely the gamma distribution with shape d and mean $d\Delta/\varepsilon$. The moment generating function is therefore

$$\mathop{\mathbb{E}}_{Z\sim\mathcal{D}^*}\left[e^{tZ}\right] = \left(1 - \frac{\Delta}{\varepsilon}t\right)^{-d}$$

for all $t < \varepsilon/\Delta$. By Markov's inequality

$$\mathop{\mathbb{P}}_{Z\sim\mathcal{D}^*}[Z \geq \alpha] \leq \frac{\mathop{\mathbb{E}}_{Z\sim\mathcal{D}^*}\left[e^{tZ}\right]}{e^{t\alpha}} = \left(1 - \frac{\Delta}{\varepsilon}t\right)^{-d} e^{-t\alpha}.$$

Setting $t = \varepsilon/\Delta - d/\alpha$ gives the required bound.

It is easy to verify that $Y \sim \mathcal{D}$ can be sampled by first sampling a radius R from a gamma distribution with shape $d+1$ and mean $(d+1)\Delta/\varepsilon$ and then sampling $Y \in [\pm R]^d$ uniformly at random. To sample R we can set $R = \frac{\Delta}{\varepsilon}\sum_{i=0}^{d} \log U_i$, where each $U_i \in (0,1]$ is uniform and independent. This gives an algorithm (in the form of an explicit circuit) to sample $\mathcal{D}$ that uses only $O(d)$ real arithmetic operations, $d+1$ logarithms, and $2d+1$ independent uniform samples from $[0,1]$.

□

4.2 Approximate Differential Privacy

Our algorithm for approximate differential privacy makes use of a powerful tool from the literature [DNR+09, HR10, DNPR10, RR10] called the sparse vector algorithm:

Theorem 4.2 (Sparse Vector). *For every $c, k \geq 1$, $\varepsilon, \delta, \alpha, \beta > 0$, and*

$$n \geq O\left(\frac{\log(k/\beta)}{\alpha\varepsilon} \min\left\{c, \sqrt{c\log(1/\delta)}\right\}\right),$$

there exists a mechanism $\mathcal{SV}$ with the following properties.

- *$\mathcal{SV}$ takes as input a database $D \in \mathcal{X}^n$ and provides answers $a_1, \cdots, a_k \in [\pm 1]$ to k (*adaptive*) linear queries $q_1, \cdots, q_k : \mathcal{X} \to [\pm 1]$.*
- *$\mathcal{SV}$ is (ε, δ)-differentially private.*
- *Assuming*
$$\left|\left\{j \in [k] : |q_j(D)| > \alpha/2\right\}\right| \leq c,$$
we have
$$\mathop{\mathbb{P}}_{\mathcal{SV}}\left[\forall j \in [k] \;\; |a_j - q_j(D)| \leq \alpha\right] \geq 1 - \beta.$$

A proof of this theorem can be found in [DR13, Theorem 3.28].[4]

Algorithm 1 Mechanism $\mathcal{M} : \{\pm 1\}^{n \times d} \to [\pm 1]^d$

Parameters: $\varepsilon, \delta > 0$.
Input: $D \in \{\pm 1\}^{n \times d}$.
Let $\sigma = 5\sqrt{d\log(1/\delta)}/\varepsilon n$ and $\alpha = 8\sqrt{\log\log d}\,\sigma$.
For $j \in [d]$, let $\tilde{a}_j = \overline{D}_j + z_j$ where $z_j \sim N(0, \sigma^2)$.
Instantiate $\mathcal{SV}$ from Theorem 4.2 with parameters $c_{\mathcal{SV}} = 2d/\log^8 d$, $k_{\mathcal{SV}} = d$, $\varepsilon_{\mathcal{SV}} = \varepsilon/2$, $\delta_{\mathcal{SV}} = \delta/2$, $\alpha_{\mathcal{SV}} = \alpha/2$, $\beta_{\mathcal{SV}} = \exp(-\log^4 d)$.
For $j \in [d]$, define $q_j : \{\pm 1\}^d \to [\pm 1]$ by $q_j(x) = (x_j - \tilde{a}_j)/2$.
Let $\hat{a}_1, \cdots, \hat{a}_d$ be the answers to $q_1, \cdots, q_d$ given by $\mathcal{SV}$.
For $j \in [d]$, let $a_j = \tilde{a}_j + 2\hat{a}_j$.
Output $a_1, \cdots, a_d$.

Proof of Theorem 1.3. The mechanism stipulated by Theorem 1.3 is given by Algorithm 1.

Firstly, we consider the privacy of $\mathcal{M}$: $\tilde{a}$ is the output of the Gaussian mechanism with parameters to ensure that it is a $(\varepsilon/2, \delta/2)$-differentially private function of D. Likewise $\hat{a}$ is the output of $\mathcal{SV}$ with parameters to ensure that it is also a $(\varepsilon/2, \delta/2)$-differentially private function of D. Since the output is $\tilde{a} + 2\hat{a}$, composition implies that $\mathcal{M}$ as a whole is (ε, δ)-differentially private, as required.

[4] Note that the algorithms in the literature are designed to sometimes output $\perp$ as an answer or halt prematurely. To modify these algorithms into the form given by Theorem 4.2 simply output 0 in these cases.

Now we must prove accuracy. Suppose that $|\hat{a}_j - q_j(D)| \leq \alpha_{\mathcal{SV}} = \alpha/2$ for all $j \in [d]$. Then

$$\begin{aligned}
|a_j - \overline{D}_j| =& |\tilde{a}_j + 2\hat{a}_j - \overline{D}_j| \\
=& |\tilde{a}_j - \overline{D}_j + 2(q_j(D) + (\hat{a}_j - q_j(D)))| \\
\leq& |\tilde{a}_j - \overline{D}_j + 2q_j(D)| + 2|\hat{a}_j - q_j(D))| \\
\leq& |\tilde{a}_j - \overline{D}_j + (\overline{D} - \tilde{a}_j)| + 2\alpha_{\mathcal{SV}} \\
=& \alpha,
\end{aligned}$$

as required. So we need only show that $|\hat{a}_j - q_j(D)| \leq \alpha_{\mathcal{SV}}$ for all $j \in [d]$, which sparse vector guarantees will happen with probability at least $1 - \beta_{\mathcal{SV}}$ as long as

$$\left|\left\{j \in [d] : |q_j(D)| > \alpha_{\mathcal{SV}}/2\right\}\right| \leq c_{\mathcal{SV}}. \tag{9}$$

Now we verify that (9) holds with high probability.

By our setting of parameters, we have $q_j(D) = -z_j/2$. This means

$$\mathbb{P}\left[|q_j(D)| > \alpha_{\mathcal{SV}}/2\right] = \mathbb{P}\left[|z_j| > \alpha/2\right] \leq e^{-\alpha^2/8\sigma^2} = \frac{1}{\log^8 d}.$$

Let $E_j \in \{0,1\}$ be the indicator of the event $|q_j(D)| > \alpha_{\mathcal{SV}}/2$. Since the z_js are independent, so are the E_js. Thus we can apply a Chernoff bound:

$$\mathbb{P}\left[\left|\left\{j \in [d] : |q_j(D)| > \alpha_{\mathcal{SV}}/2\right\}\right| > c_{\mathcal{SV}}\right] = \mathbb{P}\left[\sum_{j\in[d]} E_j > \frac{2d}{\log^8 d}\right] \leq e^{-2d/\log^{16} d}. \tag{10}$$

The failure probability of $\mathcal{M}$ is bounded by the failure probability of $\mathcal{SV}$ plus (10), which is dominated by $\beta_{\mathcal{SV}} = \exp(-\log^4 d)$.

Finally we consider the sample complexity. The accuracy is bounded by $\alpha \leq 40\sqrt{d \cdot \log(1/\delta) \cdot \log\log d}/\varepsilon n$, which rearranges to

$$n \geq \frac{40\sqrt{d \cdot \log(1/\delta) \cdot \log\log d}}{\alpha\varepsilon}.$$

Theorem 4.2 requires

$$n \geq O\left(\frac{\sqrt{c_{\mathcal{SV}} \log(1/\delta)}\log(d/\beta_{\mathcal{SV}})}{\alpha\varepsilon}\right) = O\left(\frac{\sqrt{d\log(1/\delta)}}{\alpha\varepsilon}\right)$$

for sparse vector to work, which is also satisfied. □

References

[BDMN05] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the sulq framework. In *PODS*, pages 128–138. ACM, June 13–15 2005.

[BS98] Dan Boneh and James Shaw. Collusion-secure fingerprinting for digital data. *IEEE Transactions on Information Theory*, 44(5):1897–1905, 1998.

[BUV14] Mark Bun, Jonathan Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *STOC*, pages 1–10. ACM, May 31 – June 3 2014.

[DKM+06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT*, pages 486–503. Springer, May 28–June 1 2006.

[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, pages 265–284. Springer, March 4-7 2006.

[DN03] Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *PODS*, pages 202–210. ACM, June 9-12 2003.

[DN04] Cynthia Dwork and Kobbi Nissim. Privacy-preserving datamining on vertically partitioned databases. In *CRYPTO*, pages 528–544. Springer, Aug 15–19 2004.

[DNPR10] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, STOC '10, pages 715–724, New York, NY, USA, 2010. ACM.

[DNR+09] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *STOC*, pages 381–390. ACM, May 31 - June 2 2009.

[DR13] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(34):211–407, 2013.

[HR10] Moritz Hardt and Guy Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *Proc. 51st Foundations of Computer Science (FOCS)*, pages 61–70. IEEE, 2010.

[HT10] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, STOC '10, pages 705–714, New York, NY, USA, 2010. ACM.

[KS08] Shiva Prasad Kasiviswanathan and Adam Smith. On the "semantics" of differential privacy: A bayesian formulation. *CoRR*, abs/0803.3946, 2008.

[MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '07, pages 94–103, Washington, DC, USA, 2007. IEEE Computer Society.

[RR10] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In *Proc. 42nd Symposium on Theory of Computing (STOC)*, pages 765–774. ACM, 2010.

[SU14] Thomas Steinke and Jonathan Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. *CoRR*, abs/1410.1228, 2014.

[Tar08] Gábor Tardos. Optimal probabilistic fingerprint codes. *J. ACM*, 55(2), 2008.

A Alternative Lower Bound for Pure Differential Privacy

It is known [HT10] that any ε-differentially private mechanism that answers d one-way marginals requires $n \geq \Omega(d/\varepsilon)$ samples. Our techniques yield an alternative simple proof of this fact.

Theorem A.1. *Let $M : \{\pm 1\}^{n\times d} \to [\pm 1]^d$ be a ε-differentially private mechanism. Suppose*

$$\forall D \in \{\pm 1\}^{n\times d} \quad \mathop{\mathbb{E}}_{M}\left[\left\|M(D) - \overline{D}\right\|_1\right] \leq 0.9d$$

Then $n \geq \Omega(d/\varepsilon)$.

The proof uses a special case of Hoeffding's Inequality:

Lemma A.2 (Hoeffding's Inequality)**.** *Let $X \in \{\pm 1\}^n$ be uniformly random and $a \in \mathbb{R}^n$ fixed. Then*

$$\mathop{\mathbb{P}}_{X}[\langle a, X\rangle > \lambda \|a\|_2] \leq e^{-\lambda^2/2}$$

for all $\lambda \geq 0$.

Proof of Theorem A.1. Let $x, x' \in \{\pm 1\}^d$ be independent and uniform. Let $D \in \{\pm 1\}^{n\times d}$ be n copies of x and, likewise, let $D' \in \{\pm 1\}^{n\times d}$ be n copies of x'. Let $Z = \langle M(D), x\rangle$ and $Z' = \langle M(D'), x\rangle$.

Now we give conflicting tail bounds for Z and Z', which we can relate by privacy.

By our hypothesis and Markov's inequality,

$$\begin{aligned}
\mathbb{P}[Z \leq d/20] =& \mathbb{P}[\langle M(D), x\rangle \leq 0.05d] \\
=& \mathbb{P}\left[\langle \overline{D}, x\rangle - \langle \overline{D} - M(D), x\rangle \leq 0.05d\right] \\
=& \mathbb{P}\left[\langle \overline{D} - M(D), x\rangle \geq 0.95d\right] \\
\leq& \mathbb{P}\left[\left\|\overline{D} - M(D)\right\|_1 \geq 0.95d\right] \\
\leq& \frac{\mathbb{E}\left[\left\|\overline{D} - M(D)\right\|_1\right]}{0.95d} \\
\leq& \frac{0.9}{0.95} < 0.95.
\end{aligned}$$

Since $M(D')$ is independent from x, we have

$$\forall \lambda \geq 0 \quad \mathbb{P}\left[Z' > \lambda\sqrt{d}\right] \leq \mathbb{P}\left[\langle M(D'), x\rangle > \lambda \left\|M(D')\right\|_2\right] \leq e^{-\lambda^2/2},$$

by Lemma A.2. In particular, setting $\lambda = \sqrt{d}/20$ gives $\mathbb{P}[Z' > d/20] \leq e^{-d/800}$.

Now D and D' are databases that differ in n rows, so privacy implies that

$$\mathbb{P}[M(D) \in S] \leq e^{n\varepsilon}\mathbb{P}[M(D') \in S]$$

for all S. Thus

$$\frac{1}{20} < \mathbb{P}\left[Z > \frac{d}{20}\right] = \mathbb{P}[M(D) \in S_x] \leq e^{n\varepsilon}\mathbb{P}[M(D') \in S_x] = e^{n\varepsilon}\mathbb{P}\left[Z' > \frac{d}{20}\right] \leq e^{n\varepsilon}e^{-d/800},$$

where

$$S_x = \left\{ y \in [\pm 1]^d : \langle y, x \rangle > \frac{d}{20} \right\}.$$

Rearranging $1/20 < e^{n\varepsilon} e^{-d/800}$, gives

$$n > \frac{d}{800\varepsilon} - \frac{\log(20)}{\varepsilon},$$

as required. □