

Privately Answering Counting Queries with Generalized Gaussian Mechanisms

Arun Ganesh ✉

Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, CA, USA

Jiazheng Zhao ✉

Computer Science Department, Stanford University, CA, USA

Abstract

We give the first closed-form privacy guarantees for the Generalized Gaussian mechanism (the mechanism that adds noise x to a vector with probability proportional to $\exp(-(\|x\|_p/\sigma)^p)$ for some σ, p), in the setting of answering k counting (i.e. sensitivity-1) queries about a database with (ϵ, δ) -differential privacy (in particular, with low ℓ_∞ -error). Just using Generalized Gaussian noise, we obtain a mechanism such that if the true answers to the queries are the vector d , the mechanism outputs answers \tilde{d} with the ℓ_∞ -error guarantee:

$$\mathbb{E} [\|\tilde{d} - d\|_\infty] = O\left(\frac{\sqrt{k \log \log k \log(1/\delta)}}{\epsilon}\right).$$

This matches the error bound of [18], but using a much simpler mechanism. By composing this mechanism with the sparse vector mechanism (generalizing a technique of [18]), we obtain a mechanism improving the $\sqrt{k \log \log k}$ dependence on k to $\sqrt{k \log \log \log k}$. Our main technical contribution is showing that certain powers of Generalized Gaussians, which follow a Generalized Gamma distribution, are sub-gamma.

In subsequent work, the optimal ℓ_∞ -error bound of $O(\sqrt{k \log(1/\delta)}/\epsilon)$ has been achieved by [4] and [9] independently. However, the Generalized Gaussian mechanism has some qualitative advantages over the mechanisms used in these papers which may make it of interest to both practitioners and theoreticians, both in the setting of answering counting queries and more generally.

2012 ACM Subject Classification Security and privacy → Privacy-preserving protocols

Keywords and phrases Differential privacy, counting queries, Generalized Gaussians

Digital Object Identifier 10.4230/LIPIcs.FORC.2021.1

Related Version *Previous Version*: <https://arxiv.org/abs/2010.01457>

Funding *Arun Ganesh*: Supported in part by NSF Award CCF-1535989.

Acknowledgements The inspiration for this project was a suggestion by Kunal Talwar that Generalized Gaussians could achieve the same asymptotic worst-case errors for query response as the mechanism of Steinke and Ullman. In particular, he suggested a proof sketch of a statement similar to Lemma 14 which was the basis for our proof that lemma. We are also thankful to the anonymous reviewers for multiple helpful suggestions on improving the paper.

1 Introduction

A fundamental question in data analysis is to, given a database, release answers to k numerical queries about a database d , balancing the goals of preserving the *privacy* of the individuals whose data comprises the database and preserving the *utility* of the answers to the queries. A standard formal guarantee for privacy is (ϵ, δ) -differential privacy [6, 5]. A mechanism \mathcal{M} that takes database d as input and outputs (a distribution over) answers \tilde{d} to the queries is (ϵ, δ) -differentially private if for any two databases d, d' which differ by only one individual and for any set of outcomes S , we have:



© Arun Ganesh and Jiazheng Zhao;

licensed under Creative Commons License CC-BY 4.0

2nd Symposium on Foundations of Responsible Computing (FORC 2021).

Editors: Katrina Ligett and Swati Gupta; Article No. 1; pp. 1:1–1:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$$\Pr_{\tilde{d} \sim \mathcal{M}(d)} [\tilde{d} \in S] \leq e^\epsilon \Pr_{\tilde{d} \sim \mathcal{M}(d')} [\tilde{d} \in S] + \delta. \quad (1)$$

When $\delta = 0$, this property is referred to ϵ -differential privacy. Without loss of generality, we will treat d (resp. \tilde{d}) as a k -dimensional vector corresponding to the answers to the queries (resp. the answers outputted by the mechanism). In this paper, we focus on the setting of *counting queries*, i.e. queries for which the presence of each individual in the database affects the answers by at most 1. In turn, throughout the paper we say a mechanism taking vectors in \mathbb{R}^k as input and outputting distributions over \mathbb{R}^k is (ϵ, δ) -differentially private if (1) holds for any two k -dimensional vectors d, d' such that $\|d - d'\|_\infty \leq 1$ and any subset S of \mathbb{R}^k .

To balance the goals of privacy and utility, we seek a mechanism \mathcal{M} that minimizes some objective function of the (distribution of) additive errors $\tilde{d} - d$, while satisfying (1). One natural and well-understood objective function is the ℓ_1 -error $\|\tilde{d} - d\|_1/k$, which gives the average absolute error of the answers to the queries. The well-known and simple *Laplace mechanism* [6], which outputs $\tilde{d} = d + x$ with probability proportional to $\exp(-\|x\|_1/\sigma)$ for an appropriate value of σ , achieves expected ℓ_1 -error of $O(\min\{\sqrt{k \log(1/\delta)}, k\}/\epsilon)$. A line of works on lower bounds [11, 3] culminated in a result of [18] showing this is optimal up to constants.

A less well-understood objective function is the ℓ_∞ -error $\|\tilde{d} - d\|_\infty$, which gives the maximum absolute error of the answers to the queries. The maximum absolute error is of course a more strict objective function than the average absolute error; indeed, the Laplace mechanism only achieves error $O(k \log k/\epsilon)$ and the Gaussian mechanism (which outputs $\tilde{d} = d + x$ with probability proportional to $\exp(-\|x\|_2^2/\sigma^2)$ for an appropriate value of σ) achieves error $O(\sqrt{k \log k \log(1/\delta)}/\epsilon)$. The first improvements on ℓ_∞ -error over the Laplace and Gaussian mechanisms were given by [18]¹. To summarize, the results of that paper (which prior to this paper were all the best known results) are:

- An ϵ -differentially private mechanism satisfying:

$$\Pr_{\tilde{d} \sim \mathcal{M}(d)} \left[\|\tilde{d} - d\|_\infty \geq O\left(\frac{k}{\epsilon}\right) \right] \leq e^{-\Omega(k)}, \quad (2)$$

(this matches a lower bound of [10] up to constants).

- An (ϵ, δ) -differentially private mechanism satisfying:

$$\Pr_{\tilde{d} \sim \mathcal{M}(d)} \left[\|\tilde{d} - d\|_\infty \geq O\left(\frac{\sqrt{k \log \log k \log(1/\delta)}}{\epsilon}\right) \right] \leq e^{-\log^{\Omega(1)} k}. \quad (3)$$

The mechanism achieving (3) starts by taking the Gaussian mechanism, and then uses the sparse vector mechanism to correct the entries of x with large error in a private manner.

- A lower bound showing any (ϵ, δ) -differentially private mechanism must satisfy:

$$\mathbb{E}_{\tilde{d} \sim \mathcal{M}(d)} [\|\tilde{d} - d\|_\infty] \geq \Omega\left(\frac{\sqrt{k \log(1/\delta)}}{\epsilon}\right). \quad (4)$$

¹ Their paper considers the problem setting where queries ask what fraction of n individuals satisfy some property, i.e. queries have sensitivity $1/n$ instead of 1, and the goal is to find the minimum n needed to achieve error at most α . Achieving error Δ with probability $1 - \rho$ in our setting is equivalent to needing $n \geq \Delta/\alpha$ to achieve error α with probability $1 - \rho$ in their setting.

The additional $\sqrt{\log k}$ term in the Gaussian mechanism's error bound comes from the fact that Gaussians' largest entries are roughly $\sqrt{\log k}$ times larger than their average entries. More generally, if we consider sampling x with probability proportional to $\exp(-(\|x\|_p/\sigma)^p)$ for some σ, p , the largest entry will be roughly $\log^{1/p} k$ times larger than the average entry. We refer to this distribution as the *Generalized Gaussian with shape p and scale σ* , as is it referred to in e.g. [17]. This leads to a natural question answered in this paper: What error bounds can we get by instead using *Generalized Gaussian mechanisms*?

1.1 Our Results and Techniques

Our first result is as follows:

► **Theorem 1.** *For all $1 \leq p \leq \log k$, $\epsilon \leq O(1)$, $\delta \in [2^{-O(k/p)}, 1/k]$, there exists a (ϵ, δ) -differentially private mechanism \mathcal{M} that takes in a vector $d \in \mathbb{R}^k$ and outputs a random $\tilde{d} \in \mathbb{R}^k$ such that for some sufficiently large constant c , and all $t \geq 0$:*

$$\Pr_{\tilde{d} \sim \mathcal{M}_\epsilon^p(d)} \left[\|\tilde{d} - d\|_\infty \geq \frac{ct\sqrt{kp} \log^{1/p} k \sqrt{\log(1/\delta)}}{\epsilon} \right] \leq e^{-t^p \log k}$$

In particular, this implies:

$$\mathbb{E}_{\tilde{d} \sim \mathcal{M}(d)} [\|\tilde{d} - d\|_\infty] = O \left(\frac{\sqrt{kp} \log^{1/p} k \sqrt{\log(1/\delta)}}{\epsilon} \right).$$

We also have for all $1 \leq q \leq p$:

$$\mathbb{E}_{\tilde{d} \sim \mathcal{M}(d)} \left[\frac{\|\tilde{d} - d\|_q}{k^{1/q}} \right] = O \left(\frac{\sqrt{kp} \log(1/\delta)}{\epsilon} \right).$$

We note that the lower bound on δ in Theorem 1 can easily be removed: if δ is smaller than $2^{-O(k/p)}$, we can instead use the mechanism achieving (2), which matches the error guarantees of Theorem 1 in this range of δ . The mechanism is simply to add noise from a Generalized Gaussian with shape p and an appropriate scale parameter σ . In our analysis, we arrive at the bounds $c \leq 2094$ and $\sigma \leq 262 \cdot \frac{\sqrt{kp} \log(1/\delta)}{\epsilon}$, although we did not attempt to optimize the constants in favor of a simpler analysis and presentation. We believe the multiplicative constants in both bounds can be substantially improved with a more careful analysis.

Setting $p = \Theta(\log \log k)$, this result matches the asymptotic error bound of (3). However, this result improves on (3) qualitatively. Although the mechanism achieving (3) is already not too complex, the Generalized Gaussian mechanism we use is even simpler, just adding noise from a well-known distribution. Notably, Generalized Gaussian mechanisms retain the property of the Gaussian mechanism that the noise added to each entry of d is independent (unlike the mechanism giving (3), which uses dependent noise), and that the noise has a known closed-form distribution that is easy to sample from². To the best of our knowledge, this is the first analysis giving privacy guarantees for Generalized Gaussian mechanisms besides that in [14]. Even then, [14] does not give any closed-form bounds on the value of σ needed for privacy. This analysis may be of independent interest for other applications where one would normally use the Gaussian mechanism, but may want to use a Generalized Gaussian mechanism with $p > 2$ to trade average-case error guarantees for better worst-case error guarantees.

² see e.g. https://sccn.ucsd.edu/wiki/Generalized_Gaussian_Probability_Density_Function.

We give a summary of our analysis here; the full analysis is given in Section 2. We first need to determine what value of σ causes the Generalized Gaussian mechanism to be private. Viewing the Generalized Gaussian mechanism as an instance of the exponential mechanism of [15], this reduces to deriving a tail bound on $\|x + 1\|_p^p - \|x\|_p^p$ for x sampled from the noise distribution. If p is even this is roughly equal to $p \sum_{j=1}^k x_j^{p-1}$. By a Chernoff bound on the signs of each random variable in the sum, this is roughly tail bounded by the sum of $\sqrt{k \log(1/\delta)}$ of the x_j^{p-1} random variables. These variables are distributed according to a *Generalized Gamma* distribution, which we prove is sub-gamma in Section B. This gives us the desired tail bound, and thus an upper bound on the σ needed to ensure (ϵ, δ) -differential privacy. To prove the error guarantees, we derive tail bounds on the ℓ_p -norm of x sampled from Generalized Gaussian distributions, as well as on the coordinates of points sampled from unit-radius ℓ_p -spheres, the latter of which is done by upper bounding the volume of “sphere caps” of these spheres.

Building on this result, we improve the previous best-known ℓ_∞ error for answering counting queries with (ϵ, δ) -differential privacy:

► **Theorem 2.** *For all $\epsilon \leq O(1)$, $\delta \in [2^{-O(k/\log \log \log k)}, 1/k]$, $t \in [0, O(\frac{\log k}{\log \log k})]$, there exists a (ϵ, δ) -differentially private mechanism \mathcal{M} that takes in a vector $d \in \mathbb{R}^k$ and outputs a random $\tilde{d} \in \mathbb{R}^k$ such that for a sufficiently large constant c :*

$$\Pr_{\tilde{d} \sim \mathcal{M}(d)} \left[\|\tilde{d} - d\|_\infty \geq \frac{ct \sqrt{k \log \log \log k \log(1/\delta)}}{\epsilon} \right] \leq e^{-\log^t k}.$$

In particular, if we choose e.g. $t = 2$ we get:

$$\mathbb{E}_{\tilde{d} \sim \mathcal{M}(d)} [\|\tilde{d} - d\|_\infty] = O \left(\frac{\sqrt{k \log \log \log k \log(1/\delta)}}{\epsilon} \right).$$

Again, the lower bound on δ can easily be removed using the mechanism achieving (2). We arrive at this result by improving upon Generalized Gaussian mechanisms in the same manner [18] improves upon the Gaussian mechanism: After sampling x from a Generalized Gaussian, we apply the sparse vector mechanism to x to get \tilde{x} which satisfies $\|x - \tilde{x}\|_\infty \ll \|x\|_\infty$. We then just output $\tilde{d} = d + x - \tilde{x}$. The full analysis is given in Section 3. Similarly to [18], the major technical component is showing that at least $k/\log^{\Omega(1)} k$ entries of x are small with high probability, which we do using the tail bounds derived in Section 2. This is necessary for the sparse vector mechanism to satisfy that $\|x - \tilde{x}\|_\infty$ is, roughly speaking, the $(k/\log^{\Omega(1)} k)$ -th largest entry of x rather than the largest entry with high probability.

1.2 Subsequent Work and Comparisons

Following our work, [4] and [9] independently gave mechanisms with optimal expected ℓ_∞ -error $O(\sqrt{k \log(1/\delta)}/\epsilon)$, quantitatively improving our results. Since in practice $\sqrt{\log \log k}$ is unlikely to be much larger than the constants hidden by the asymptotic notation (e.g., using the natural log, $\sqrt{\log \log k} = 2$ for $k \approx 5 \cdot 10^{23}$), the qualitative differences between our results and these two results make our results still of interest to e.g. practitioners. Theorem 1 is our qualitatively more appealing result, and so we highlight the differences with that result in particular. Again, we note that while the explicit constant in Theorem 1 is likely too large to be of practical interest, we believe this constant can be substantially improved with a more refined analysis, hopefully making the mechanism practical.

The result of [4] remarkably uses a bounded noise distribution, and in turn the *maximum* ℓ_∞ -error rather than just the average ℓ_∞ -error of their mechanism is bounded, in contrast with Generalized Gaussian mechanisms whose maximum ℓ_∞ -error is unbounded. However, a bounded noise distribution cannot e.g. satisfy group differential privacy for all group sizes simultaneously, whereas Generalized Gaussian mechanisms can. Also, while both results simply add noise, Generalized Gaussians are more well-studied than the noise distribution of [4] and can be sampled by simply powering and rescaling samples from Gamma random variables, which should make them easier to implement in practice.

The result of [9] at a high level adds noise and then repeatedly applies the sparse vector mechanism to correct entries with large noise, in contrast to just adding noise. In addition, their result uses arguably even simpler sampling primitives than ours (it only needs to sample Laplace distributions and permutations of lists), but their overall mechanism needs a somewhat more involved iterative approach rather than a one-shot sample. Finally, as presented the resulting noise distribution from their overall mechanism does not have e.g. a closed-form or independent entries which may be desirable.

1.3 Preliminaries

For completeness, we restate the noise distribution of interest here:

► **Definition 3.** *The (multivariate) **Generalized Gaussian distribution with shape p and scale σ** denoted $GGauss(p, \sigma)$, is the distribution over $x \in \mathbb{R}^k$ with probability distribution function (pdf) proportional to $\exp(-(\|x\|_p/\sigma)^p)$.*

1.3.1 Sub-Gamma Random Variables

The following facts about sub-gamma random variables will be useful in our analysis:

► **Definition 4.** *A random variable X is **sub-gamma to the right** with variance v and scale c if:*

$$\forall \lambda \in (0, 1/c) : \mathbb{E}[\exp(\lambda(X - \mathbb{E}[X]))] \leq \exp\left(\frac{\lambda^2 v}{2(1 - c\lambda)}\right).$$

*Here, we use the convention $1/c = \infty$ if $c = 0$. We denote the class of such random variables $\Gamma^+(v, c)$. Similarly, a random variable X is **sub-gamma to the left** with variance v and scale c , if $-X \in \Gamma^+(v, c)$, i.e.:*

$$\forall \lambda \in (0, 1/c) : \mathbb{E}[\exp(\lambda(\mathbb{E}[X] - X))] \leq \exp\left(\frac{\lambda^2 v}{2(1 - c\lambda)}\right).$$

We denote the class of such random variables $\Gamma^-(v, c)$.

We refer the reader to [1] for a textbook reference for this definition and proofs of the following facts.

► **Fact 5.** *If for $i \in [n]$ we have a random variable $X_i \in \Gamma^+(v_i, c_i)$, then $X = \sum_{i \in [n]} X_i$ satisfies $X \in \Gamma^+(\sum_{i \in [n]} v_i, \max_{i \in [n]} c_i)$ (and the same relation holds for $\Gamma^-(v, c)$).*

► **Lemma 6.** *If $X \in \Gamma^+(v, c)$ then for all $t > 0$:*

$$\Pr[X > \mathbb{E}[X] + \sqrt{2vt} + ct] \leq e^{-t}.$$

Similarly, if $X \in \Gamma^-(v, c)$ then for all $t > 0$:

$$\Pr[X < \mathbb{E}[X] - \sqrt{2vt} - ct] \leq e^{-t}.$$

► **Fact 7.** Let $X \sim \text{Gamma}(a)$, i.e. X has pdf satisfying:

$$p(x) \propto x^{a-1} e^{-x}.$$

Then X satisfies $X \in \Gamma^+(a, 1)$ and $X \in \Gamma^-(a, 0)$.

1.3.2 Other Probability Facts

We will use the following standard fact to relate distributions of variables to the distributions of their powers:

► **Fact 8 (Change of Variables for Powers).** Let X be distributed over $(0, \infty)$ with pdf proportional to $f(x)$. Let Y be the random variable X^c for $c > 0$. Then Y has pdf proportional to $y^{\frac{1}{c}-1} f(y^{\frac{1}{c}})$.

Finally, we'll use the following standard tail bounds:

► **Lemma 9 (Laplace Tail Bound).** Let X be a Laplace random variable with scale b , $\text{Lap}(b)$. That is, X has pdf proportional to $\exp(-|x|/b)$. Then we have $\Pr[|x| \geq tb] \leq e^{-t}$.

► **Lemma 10 (Chernoff Bound).** Let X_1, X_2, \dots, X_k be independent Bernoulli random variables. Let $\mu = \mathbb{E} \left[\sum_{i \in [k]} X_i \right]$. Then for $t \in [0, 1]$, we have:

$$\Pr \left[\sum_{i \in [k]} X_i \geq (1+t)\mu \right] \leq \exp \left(-\frac{t^2 \mu}{3} \right).$$

2 Generalized Gaussian Mechanisms

In this section, we analyze the Generalized Gaussian mechanism that given database d , samples $x \sim \text{GGauss}(p, \sigma)$ and outputs $\tilde{d} = d + x$. We denote this mechanism \mathcal{M}_σ^p . When $p = 1$ this is the Laplace mechanism, and when $p = 2$ this is the Gaussian mechanism.

2.1 Privacy Guarantees

We first determine what σ is needed to make this mechanism private. We start with the following lemma, which gives a tail bound on the change in the “utility” function $\|\tilde{d} - d\|_p^p$ if d changes by $\Delta \in [-1, 1]^k$:

► **Lemma 11.** Let $x \in \mathbb{R}^k$ be sampled from $\text{GGauss}(p, \sigma)$. Then for $4 \leq p \leq \log k$ that is an even integer, $\delta \in [2^{-O(k/p)}, 1/k]$, and any $\Delta \in [-1, 1]^k$ we have with probability $1 - \delta$:

$$\|x - \Delta\|_p^p - \|x\|_p^p \leq 32pk^{1/p-1/2} \sqrt{p \log(1/\delta)} \|x\|_p^{p-1} + 2k^{\frac{p}{2}} p^2.$$

We remark that the requirement that p be an even integer can be dropped by generalizing the proofs in this section appropriately. However, we can reduce proving Theorem 1 for all p to proving it for only even p by rounding p up to the nearest even integer (at the loss of a multiplicative constant of at most $\sqrt{2}$), and only considering even p simplifies the presentation. So, we stick to considering only even p .

Proof. By symmetry of $\text{GGauss}(p, \sigma)$ we can assume Δ has all negative entries. Then we have:

$$\begin{aligned} \|x - \Delta\|_p^p - \|x\|_p^p &= \sum_{i=1}^k ((x_i - \Delta_i)^p - x_i^p) \\ &= \sum_{i=1}^k \int_{x_i}^{x_i - \Delta_i} p y^{p-1} dy \leq \sum_{i=1}^k \int_{x_i}^{x_i - \Delta_i} p (x_i - \Delta_i)^{p-1} dy \leq p \sum_{i=1}^k (x_i - \Delta_i)^{p-1} \leq p \sum_{i=1}^k (x_i + 1)^{p-1}. \end{aligned}$$

We want to replace the terms $(x_i + 1)^{p-1}$ with terms x_i^{p-1} since the latter's distribution is more easily analyzed. To do so, we use the following observation:

► **Fact 12.** *If $p \leq \sqrt{k}/2$:*

- *If $x_i > \sqrt{k}$, then we have $(x_i + 1)^{p-1} \leq \left(1 + \frac{1}{\sqrt{k}}\right)^{p-1} x_i^{p-1} \leq \left(1 + \frac{2p}{\sqrt{k}}\right) x_i^{p-1}$.*
- *If $|x_i| \leq \sqrt{k}$, then we have $(x_i + 1)^{p-1} - x_i^{p-1} \leq (\sqrt{k} + 1)^{p-1} - \sqrt{k}^{p-1} \leq 2k^{\frac{p}{2}-1}p$.*
- *If $x_i < -\sqrt{k}$, then we have $(x_i + 1)^{p-1} \leq \left(1 - \frac{1}{\sqrt{k}}\right)^{p-1} x_i^{p-1} \leq \left(1 - \frac{2p}{\sqrt{k}}\right) x_i^{p-1}$.*

Fact 12 gives:

$$\sum_{i=1}^k (x_i + 1)^{p-1} \leq \left(1 - \frac{2p}{\sqrt{k}}\right) \sum_{i:x_i < 0} x_i^{p-1} + \left(1 + \frac{2p}{\sqrt{k}}\right) \sum_{i:x_i \geq 0} x_i^{p-1} + 2k^{\frac{p}{2}}p.$$

It now suffices to show that:

$$-\left(1 - \frac{2p}{\sqrt{k}}\right) \sum_{i:x_i < 0} |x_i|^{p-1} + \left(1 + \frac{2p}{\sqrt{k}}\right) \sum_{i:x_i \geq 0} |x_i|^{p-1} \leq 32k^{1/p-1/2} \sqrt{p \log(1/\delta)} \|x\|_p^{p-1}, \quad (5)$$

with probability at least $1 - \delta$. Note that each x_i is sampled independently with probability proportional to $\exp(-(|x_i|/\sigma)^p)$. Since multiplying x by a constant rescales both sides of (5) by the same multiplicative factor, it suffices to show (5) when each x_i is independently sampled with probability proportional to $\exp(-|x_i|^p)$, i.e. when $\sigma = 1$. By change of variables, $y_i = |x_i|^{p-1}$ is sampled from the distribution with pdf proportional to $y_i^{\frac{1}{p-1}-1} \exp(-y_i^{\frac{p}{p-1}})$. This is the Generalized Gamma random variable with parameters $(\frac{1}{p-1}, \frac{p}{p-1})$, which we denote $GGamma(\frac{1}{p-1}, \frac{p}{p-1})$. We show the following property of this random variable in Appendix B:

► **Lemma 13.** *For any $p \geq 4$, let Y be the random variable $GGamma(\frac{1}{p-1}, \frac{p}{p-1})$, let $\mu = \mathbb{E}[Y]$. Then $\mu \in [1/p, 1.2/p]$, $Y \in \Gamma^+(\mu, 1)$, and $Y \in \Gamma^-(\mu, 3/2)$.*

Let k' be the number of positive coordinates in x . A Chernoff bound gives that $k' \leq \frac{k}{2} + 3\sqrt{k \log \frac{1}{\delta}}$ with probability $1 - \delta/3$. By Lemma 13 and Fact 5 $\sum_{i:x_i < 0} |x_i|^{p-1}$ is in $\Gamma^-(k - k')\mu, 3/2)$ and $\sum_{i:x_i \geq 0} |x_i|^{p-1}$ is in $\Gamma^+(k'\mu, 1)$ for μ as defined in Lemma 13. We now apply Lemma 6 with $t = \log(6/\delta)$ to each sum. Since $\delta \geq 2^{-O(k/\sqrt{p})}$, $\log(6/\delta) = O(\sqrt{k \log(1/\delta)/p})$, i.e. we are still in the range of δ for which the square-root term in the tail bound of Lemma 6 is the linear term ct . So Lemma 6 gives that:

$$\Pr \left[\sum_{i:x_i < 0} |x_i|^{p-1} < (k - k')\mu - 2\sqrt{2k\mu \log(1/\delta)} \right] \leq \delta/6,$$

$$\Pr \left[\sum_{i:x_i \geq 0} |x_i|^{p-1} > k'\mu + 2\sqrt{2k\mu \log(1/\delta)} \right] \leq \delta/6.$$

Combined with the Chernoff bound, this gives that with probability $1 - 2\delta/3$:

$$\begin{aligned}
& - \left(1 - \frac{2p}{\sqrt{k}}\right) \sum_{i:x_i < 0} |x_i|^{p-1} + \left(1 + \frac{2p}{\sqrt{k}}\right) \sum_{i:x_i \geq 0} |x_i|^{p-1} \\
\leq & - \left(1 - \frac{2p}{\sqrt{k}}\right) \left((k - k')\mu - (2\sqrt{2})\sqrt{k\mu \log(1/\delta)}\right) \\
& + \left(1 + \frac{2p}{\sqrt{k}}\right) \left(k'\mu + (2\sqrt{2})\sqrt{k\mu \log(1/\delta)}\right) \\
\leq & (2k' - k)\mu + (2\sqrt{k}p)\mu + (4\sqrt{2})\sqrt{k\mu \log(1/\delta)} \\
\leq & 6\mu\sqrt{k \log(1/\delta)} + (2\sqrt{k}p)\mu + (5\sqrt{2})\mu\sqrt{kp \log(1/\delta)} \\
\leq & 16k\mu \cdot \frac{\sqrt{p \log(1/\delta)}}{\sqrt{k}}.
\end{aligned} \tag{6}$$

In the last step, we use that $p \leq \log k \leq \log(1/\delta)$ for the range of p, δ we consider. On the other hand, by Fact 5 $\sum_{i \in [k]} |x_i|^{p-1} = \|x\|_{p-1}^{p-1}$ is sampled from a random variable in $\Gamma^-(k\mu, 3/2)$ and thus by Lemma 13 and Lemma 6 is at least $k\mu/2$ with probability at least $1 - \delta/3$, i.e. $k\mu \leq 2\|x\|_{p-1}^{p-1}$ with probability at least $1 - \delta/3$. Combined with (7) by a union bound we get with probability $1 - \delta$:

$$- \left(1 - \frac{2p}{\sqrt{k}}\right) \sum_{i:x_i < 0} |x_i|^{p-1} + \left(1 + \frac{2p}{\sqrt{k}}\right) \sum_{i:x_i \geq 0} |x_i|^{p-1} \leq 32 \frac{\sqrt{p \log(1/\delta)}}{\sqrt{k}} \cdot \|x\|_{p-1}^{p-1}$$

Finally, by the Cauchy-Schwarz inequality for any $a \leq b$ and k -dimensional x we have $\|x\|_a \leq k^{1/a-1/b} \|x\|_b$. So, $\|x\|_{p-1}^{p-1} \leq k^{1/p} \|x\|_p^{p-1}$, giving (5) with probability $1 - \delta$ as desired. \blacktriangleleft

Given Lemma 11, determining the value of σ that makes \mathcal{M}_σ^p private is fairly straightforward:

► **Lemma 14.** *Let \mathcal{M}_σ^p be the mechanism such that $\mathcal{M}_\sigma^p(d)$ samples $x \in \mathbb{R}^k$ from $x \sim \text{GGauss}(p, \sigma)$ and outputs $\tilde{d} = d + x$. For $4 \leq p \leq \log k$ that is an even integer, $\epsilon \leq O(1)$, $\delta \in [2^{-O(k/p)}, 1/k]$, and*

$$\sigma = \Theta \left(\frac{\sqrt{kp \log(1/\delta)}}{\epsilon} \right),$$

\mathcal{M}_σ^p is (ϵ, δ) -differentially private.

Proof. It suffices to show that for any vector Δ in $[-1, 1]^k$:

$$\Pr_{\tilde{d} \sim \mathcal{M}_\sigma^p(d)} \left[\log \left(\frac{\Pr[\mathcal{M}_\sigma^p(d) = \tilde{d}]}{\Pr[\mathcal{M}_\sigma^p(d + \Delta) = \tilde{d}]} \right) \leq \epsilon \right] = \Pr_{\tilde{d} \sim \mathcal{M}_\sigma^p(d)} \left[\frac{\|x - \Delta\|_p^p - \|x\|_p^p}{\sigma^p} \leq \epsilon \right] \geq 1 - \delta.$$

Here, we abuse notation by letting \Pr also denote a likelihood function. By Lemma 11 we now have with probability $1 - \delta/2$ for a sufficiently large constant c :

$$\|x - \Delta\|_p^p - \|x\|_p^p \leq 64pk^{1/p-1/2} \sqrt{p \log(1/\delta)} \|x\|_p^{p-1} + 2p^2 k^{\frac{p}{2}}.$$

The pdf of the rescaled norm $r = \|x\|_p/\sigma$ is proportional to $r^{k-1} \exp(-r^p)$ over $(0, \infty)$ (the r^{k-1} appears because the $(k-1)$ -dimensional surface area of the ℓ_p -sphere of radius r is proportional to r^{k-1}). Letting R denote r^p , the pdf of R is proportional to $R^{\frac{k}{p}-1} \exp(-R)$

by change of variables, i.e. R is the random variable $\text{Gamma}(\frac{k}{p})$. Then by the Gamma tail bound, with probability at least $1 - e^{-.001k/p} > 1 - \delta/2$, R is contained in $[\frac{k}{2p}, \frac{2k}{p}]$, so $\|x\|_p$ is contained in $[\sigma \left(\frac{k}{2p}\right)^{1/p}, \sigma \left(\frac{2k}{p}\right)^{1/p}]$. Then by a union bound, with probability $1 - \delta$:

$$\frac{\|x - \Delta\|_p^p - \|x\|_p^p}{\sigma^p} \leq \frac{128p^{1/p} \sqrt{kp \log(1/\delta)}}{\sigma} + \frac{4p^2 k^{\frac{p}{2}}}{\sigma^p}.$$

Noting that $n^{1/n}$ is contained within $[1, e^{1/e}]$ for all $n \geq 1$, letting

$$\sigma = 185 \cdot \frac{\sqrt{kp \log(1/\delta)}}{\epsilon},$$

we get that $\frac{\|x - \Delta\|_p^p - \|x\|_p^p}{\sigma^p} \leq \epsilon$ with probability $1 - \delta$ as desired. \blacktriangleleft

2.2 Error Guarantees

In this section, we analyze the ℓ_∞ error of \mathcal{M}_σ^p , for a given choice of δ in the range specified in Lemma 14. We give an expected error bound, and also a tail bound on the error. The error analysis follows almost immediately from the following lemma, which bounds the fraction of a sphere cap's volume with a large first coordinate:

► **Lemma 15.** *Let x be chosen uniformly at random from a k -dimensional ℓ_p -sphere with arbitrary radius, i.e. the set of points with $\|x\|_p = R$ for some R , for $p \geq 1$. Then we have:*

$$\Pr[|x_1| \geq r \|x\|_p] \leq (1 - r^p)^{(k-1)/p} \leq \exp\left(-\frac{(k-1)r^p}{p}\right).$$

This lemma or one providing a similar bound likely already exists in the literature, but we are unaware of a reference for it. So, for completeness we give the full proof at the end of the section.

► **Corollary 16.** *Let x be chosen uniformly at random from a k -dimensional ℓ_p -sphere with arbitrary radius for $p \geq 1$. Then we have:*

$$\Pr[\|x\|_\infty \geq r \|x\|_p] \leq k \cdot \exp\left(-\frac{(k-1)r^p}{p}\right).$$

Proof. This follows from Lemma 15 and a union bound over all k coordinates (which have identical marginal distributions). \blacktriangleleft

Combining this corollary with Lemma 14, it is fairly straightforward to prove our first main result:

► **Theorem 17.** *Let \mathcal{M}_σ^p be the mechanism such that $\mathcal{M}_\sigma^p(d)$ samples $x \in \mathbb{R}^k$ from $\text{GGauss}(p, \sigma)$ and outputs $\tilde{d} = d + x$. For $4 \leq p \leq \log k$ that is an even integer, For $\epsilon \leq O(1)$, $\delta \in [2^{-O(k/p)}, 1/k]$, and*

$$\sigma = 185 \cdot \frac{\sqrt{kp \log(1/\delta)}}{\epsilon},$$

\mathcal{M}_σ^p is (ϵ, δ) -differentially private and for some sufficiently large constant c , and all $t \geq 0$:

$$\Pr_{\tilde{d} \sim \mathcal{M}_\sigma^p(d)} \left[\|\tilde{d} - d\|_\infty \geq 1480t \cdot \frac{\sqrt{kp \log^{1/p} k \sqrt{\log(1/\delta)}}}{\epsilon} \right] \leq e^{-t^p \log k} + e^{-.001k/p}$$

Proof. The privacy guarantee follows from Lemma 14.

For the tail bound, if $\|\tilde{d} - d\|_\infty > 1480t \cdot \frac{\sqrt{k} \log^{1/p} k \sqrt{p \log(1/\delta)}}{\epsilon}$ we have either $\|x\|_p \geq 370 \cdot \frac{k^{1/2+1/p} \sqrt{p \log(1/\delta)}}{\epsilon}$ or $\|x\|_\infty > \frac{4t \log^{1/p} k}{k^{1/p}} \|x\|_p$. Recall that $(\|x\|_p/\sigma)^p$ is distributed according to a $\text{Gamma}(\frac{k}{p})$ random variable, and thus by a Gamma tail bound exceeds $2k/p$ with probability at most $e^{-.001k/p}$. In turn, $\|x\|_p \geq 370 \cdot \frac{k^{1/2+1/p} \sqrt{p \log(1/\delta)}}{\epsilon} \geq \left(\frac{2k}{p}\right)^{1/p} \sigma$ with at most this probability. Then it follows by setting $r = \frac{4t \log^{1/p} k}{k^{1/p}}$ in Corollary 16 and a union bound that:

$$\begin{aligned} \Pr \left[\|\tilde{d} - d\|_\infty \geq 1480t \cdot \frac{\sqrt{k} \log^{1/p} k \sqrt{p \log(1/\delta)}}{\epsilon} \right] &\leq \Pr \left[\|x\|_\infty \geq \frac{4t \log^{1/p} k}{k^{1/p}} \|x\|_p \right] \\ + e^{-.001k/p} &\leq \exp \left(-\frac{(k-1)4^p t^p \log k}{kp} \right) + e^{-.001k/p} \leq e^{-t^p \log k} + e^{-.001k/p}. \quad \blacktriangleleft \end{aligned}$$

This proves Theorem 1, up to some details which we defer to Section A.

2.3 Proof of Lemma 15

To prove this lemma we'll need the following lemma about convex bodies.

► **Lemma 18.** *Let $A \subseteq B \subset \mathbb{R}^k$ be two compact convex bodies with A contained in B , and A', B' be their respective boundaries. Then $\text{Vol}_{k-1}(A') \leq \text{Vol}_{k-1}(B')$, where Vol_{k-1} denotes the $(k-1)$ -dimensional volume.*

Proof. For any compact convex body S and its boundary S' , the $(k-1)$ -dimensional volume of S' satisfies:

$$\text{Vol}_{k-1}(S') \propto \int_{\mathbb{S}^k} \text{Vol}_{k-1}(\pi_{\theta^\top} S) d\theta,$$


Where \mathbb{S}^k is the k -dimensional unit sphere and $\pi_{\theta^\top} S$ is the orthogonal projection of S onto the subspace of \mathbb{R}^k orthogonal to θ (see e.g. Section 5.5 of [13] for a proof of this fact). Since $A \subseteq B$ it follows that for all θ we have $\text{Vol}_{k-1}(\pi_{\theta^\top} A) \leq \text{Vol}_{k-1}(\pi_{\theta^\top} B)$ and so $\text{Vol}_{k-1}(A') \leq \text{Vol}_{k-1}(B')$. ◀

The idea behind the proof of Lemma 15 is to show that the region of the ℓ_p -ball with large positive first coordinate is contained within a smaller ℓ_p -ball, and then apply Lemma 18:

Proof of Lemma 15. By rescaling, we can assume $\|x\|_p = 1$ and instead show:

$$\begin{aligned} \Pr[|x_1| \geq r] &\leq (1 - r^p)^{(k-1)/p} \\ \Pr[|x_1| \geq r] &= \frac{\text{Vol}_{k-1}(\{x : |x_1| \geq r, \|x\|_p = 1\})}{\text{Vol}_{k-1}(x : \|x\|_p = 1)} = \frac{\text{Vol}_{k-1}(\{x : x_1 \geq r, \|x\|_p = 1\})}{\text{Vol}_{k-1}(\{x : x_1 \geq 0, \|x\|_p = 1\})}, \end{aligned}$$

Where Vol_{k-1} denotes the $(k-1)$ -dimensional volume. To bound this ratio, let v be the vector $(r, 0, 0, \dots, 0)$, and consider the (compact, convex) body $B_1 = \{x : x_1 \geq r, \|x - v\|_p \leq (1 - r^p)^{1/p}\}$. We have $r^p + (v - r)^p \leq v^p$ for $0 \leq r \leq v$, so B_1 contains the (also compact, convex) body $B_2 = \{x : x_1 \geq r, \|x\|_p \leq 1\}$. Then by Lemma 18 the $(k-1)$ -dimensional surface area of B_1 is larger than that of B_2 . The boundary of B_1 is the union of the bodies $B_{1,a} := \{x : x_1 = r, \|x - v\|_p \leq (1 - r^p)^{1/p}\}$ and $B_{1,b} := \{x : x_1 \geq r, \|x - v\|_p = (1 - r^p)^{1/p}\}$, whose intersection has $(k-1)$ -dimensional volume 0. Similarly, the boundary of B_2 is the



■ **Figure 1** A picture of the bodies in the proof of Lemma 15 for $p = 2, k = 2$. B_2 has stripes that are the same color as $B_1 \setminus B_2$ to emphasize that B_1 contains B_2 .

union of the bodies $B_{2,a} := \{x : x_1 = r, \|x\|_p \leq 1\}$ and $B_{2,b} := \{x : x_1 \geq r, \|x\|_p = 1\}$, whose intersection has $(k-1)$ -dimensional volume 0. See Figure 1 for an example of a picture of all of these bodies.

Nothing that $B_{1,a} = B_{2,a}$, we conclude that $\text{Vol}_{k-1}(B_{1,b}) \geq \text{Vol}_{k-1}(B_{2,b})$. Now we have:

$$\frac{\text{Vol}_{k-1}(\{x : x_1 \geq r, \|x\|_p = 1\})}{\text{Vol}_{k-1}(\{x : x_1 \geq 0, \|x\|_p = 1\})} \leq \frac{\text{Vol}_{k-1}(\{x : x_1 \geq r, \|x - v\|_p = (1 - r^p)^{1/p}\})}{\text{Vol}_{k-1}(\{x : x_1 \geq 0, \|x\|_p = 1\})}.$$

The body in the numerator of the final expression is the body in the denominator, but shifted by v and rescaled by $(1 - r^p)^{1/p}$ in every dimension. So, the final ratio is at most $(1 - r^p)^{(k-1)/p}$. ◀

3 Composition with Sparse Vector

In this section, we generalize the mechanism of [18], which is a composition of the Gaussian mechanism and sparse vector mechanism of [7], by analyzing a composition of \mathcal{M}_σ^p and the sparse vector mechanism instead³. The guarantees given by sparse vector can be given in the following form that we will use:

▶ **Theorem 19** (Sparse Vector). *For every $k \geq 1, c_{SV} \leq k, \epsilon_{SV}, \delta_{SV}, \beta_{SV} > 0$, and*

$$\alpha_{SV} \geq O\left(\frac{\sqrt{c_{SV} \log(1/\delta_{SV}) \log(k/\beta_{SV})}}{\epsilon_{SV}}\right),$$

there exists a mechanism SV that takes as input $d \in \mathbb{R}^k$ and outputs $\tilde{d} \in \mathbb{R}^k$ such that:

- *SV is $(\epsilon_{SV}, \delta_{SV})$ -differentially private.*
- *If at most c_{SV} entries of d have absolute value strictly greater than $\alpha_{SV}/2$, then:*

$$\Pr_{\tilde{d} \sim SV(d)} [\|\tilde{d} - d\|_\infty \geq \alpha_{SV}] \leq \beta_{SV}.$$

- *Regardless of the value of d we have for all $t \geq 0$:*

$$\Pr_{\tilde{d} \sim SV(d)} [\|\tilde{d} - d\| \geq \max\{\|d\|_\infty, t\sqrt{k \log(1/\delta_{SV})/\epsilon_{SV}}\}] \leq ke^{-\Omega(t)}.$$

³ Unlike its preprint, the journal version of [18] uses a slightly different mechanism based on the exponential mechanism in place of the sparse vector mechanism. A similar change can likely be made to the mechanism given in this section; we stick to using the sparse vector mechanism for a slightly simpler proof.

1:12 Privately Answering Counting Queries with Generalized Gaussian Mechanisms

The proof is deferred to Section A. We now prove Theorem 20, from which Theorem 2 follows up to some minor details (see Section A):

► **Theorem 20.** *For any $4 \leq p \leq \log k$ that is an even integer, $\epsilon \leq O(1)$, $\delta \in [2^{-O(k/p)}, 1/k]$, and $t \in [0, O(\frac{\log k}{\log \log k})]$, there exists a (ϵ, δ) -differentially private mechanism \mathcal{M} that takes in a vector $d \in \mathbb{R}^k$ and outputs a random $\tilde{d} \in \mathbb{R}^k$ such that for a sufficiently large constant c :*

$$\Pr_{\tilde{d} \sim \mathcal{M}(d)} \left[\|\tilde{d} - d\|_\infty \geq \frac{ct\sqrt{kp \log(1/\delta)}(\log \log k)^{1/p}}{\epsilon} \right] \leq e^{-\log^t k}.$$

Proof. The mechanism is as follows: We sample $x \sim GGauss(p, \sigma)$ for

$$\sigma = \Theta \left(\frac{\sqrt{kp \log(1/\delta)}}{\epsilon} \right),$$

If $\|x\|_p^p > 2k\sigma^p/p$, we output d . Otherwise, we instantiate SV from Theorem 19 with parameters:

$$\alpha_{SV} = 12t(\log \log k)^{1/p}\sigma \leq \frac{ct\sqrt{kp \log(1/\delta)}(\log \log k)^{1/p}}{\epsilon}, \quad c_{SV} = 4k/\log^{2+2t} k,$$

$$\epsilon_{SV} = \epsilon/2, \quad \delta_{SV} = \delta/3, \quad \beta_{SV} = \exp(-\log^t k)/2.$$

We input x to SV to sample \hat{x} , and then output $\tilde{d} = d + x - \hat{x}$.

First, note that:

$$\frac{\sqrt{c_{SV} \log(1/\delta_{SV})} \log(k/\beta_{SV})}{\epsilon_{SV}} \leq \frac{\sqrt{\frac{16k}{\log^{2+2t} k} \log(1/\delta)} (\log k + \log^t k)}{\epsilon} \leq \frac{4\sqrt{k \log(1/\delta)}}{\epsilon},$$

i.e. α satisfies the requirements of Theorem 19 as long as the constant hidden in the $\Theta(\cdot)$ notation in the choice of σ is sufficiently large.

To analyze the privacy guarantee, this is the composition of:

- The mechanism of Theorem 17, which if the constant hidden in the $\Theta(\cdot)$ in the expression for σ is sufficiently large, is $(\epsilon/2, \delta/3)$ -differentially private.
- The SV mechanism of Theorem 19, with parameters set so it is $(\epsilon/2, \delta/3)$ -differentially private.
- The event that $\|x\|_p^p > 2k\sigma^p/p$, causing us to release the database, which we recall from the Proof of Theorem 17 happens with probability at most $2^{-\Omega(k/p)} \leq \delta/3$.

By composition, we get that the mechanism is (ϵ, δ) -differentially private as desired.

To show the tail bound on ℓ_∞ -error: If $\|x\|_p^p > 2k\sigma^p/p$, then we have $\tilde{d} = d$, so trivially the tail bound is satisfied. So, it suffices to show that conditional on $\|x\|_p^p \leq 2k\sigma^p/p$ occurring, we have the tail bound. By a union bound, the guarantees of Theorem 19 give that $\|\tilde{d} - d\|_\infty = \|x - \hat{x}\|_\infty \leq \alpha_{SV}$ (i.e the tail bound is satisfied) if at most $4k/\log^{2+2t} k$ entries of x have absolute value greater than $\alpha_{SV}/2$ with probability less than, say, $e^{-2\log^t k}$. Using $r = 3t \frac{(\log \log k)^{1/p}}{k^{1/p}}$ in Lemma 15 and a union bound with the $1 - \delta/3$ probability event that $\|x\|_p \leq (2k/p)^{1/p}\sigma$, for each coordinate x_i of x we have:

$$|x_i| \geq \alpha_{SV}/2 = 6t(\log \log k)^{1/p}\sigma = 2rk^{1/p}\sigma \geq r\|x\|_p,$$

with probability at most $\frac{1}{\log^{2+2t} k} + 2^{-\Omega(k/p)} \leq \frac{2}{\log^{2+2t} k}$. Since we sample x with probability proportional to $\exp(-\sum_{i \in [k]} |x_i|^p/\sigma^p)$, each coordinate's distribution is independent, so using a Chernoff bound we conclude that with probability $e^{-\Omega(k/\log^{2+2t} k)} \leq e^{-2\log^t k}$ at most $4k/\log^{2+2t} k$ coordinates have absolute value greater than α_{SV} as desired. ◀

4 Future Directions

As mentioned before, we did not attempt to optimize the constant multiplier in Theorem 1, and our resulting constant is likely too large to be practical. Since the Generalized Gaussian generalizes the Laplace and Gaussian mechanisms, which have good multiplicative constants in practice, we expect that a more careful analysis of the Generalized Gaussian will also lead to a error bound that is practical.

Another question concerns stronger measures of privacy than (ϵ, δ) -DP, including Rényi-DP [16] and zero-concentrated-DP [2]. To show the Generalized Gaussian mechanism satisfies these notions of privacy requires one to bound a moment generating function of the privacy loss $\frac{\|x-\Delta\|_p^p - \|x\|_p^p}{\sigma^p}$, which in some sense requires the privacy loss to be subexponential. Roughly speaking, our analysis shows with probability at least $1 - \delta$, the privacy loss lies in an interval in which it behaves as a subgaussian random variable. However, past this interval, our analysis fails to show it even behaves subexponentially. This is because our use of the gamma tail bound of Lemma 6 weakens at two points in the regime where $\delta < 2^{-k/p}$. The first is that the final expression in (7) has a dependence on δ of $\log(1/\delta)$ instead of $\sqrt{\log(1/\delta)}$ when $\delta < 2^{-k/p}$, since the linear term ct in Lemma 6 begins to dominate the error. The second is that, roughly speaking, we use the gamma tail bound to show that $\|x\|_p^p$ deviates from its expectation of k/p by at most $\sqrt{k \log(1/\delta)/p}$ with probability $1 - \delta$. When $\delta \geq 2^{-k/p}$, this lets us treat $\|x\|_p^p$ as always being within a constant factor of its expectation in our analysis. However, when δ is small enough, the term $\sqrt{k \log(1/\delta)/p}$ becomes much larger than the term k/p , and so we can only bound $\|x\|_p^p$'s deviation from its expectation by an expression with $\sqrt{\log(1/\delta)}$ dependence on δ .

Our final tail bound on the privacy loss is effectively a product of the tail bound of Lemma 11 and the tail bound on $\|x\|_p^{p-1}$, and so it shows concentration that is worse than sub-exponential in the small δ regime, which is insufficient for proving these stronger notions of privacy. We believe this is a function of our analysis rather than of the Generalized Gaussian mechanism, but do not know of an alternate analysis that confirms this belief. Determining whether Generalized Gaussian mechanisms can satisfy stronger notions of privacy for larger values of p is an interesting open direction.

References

- 1 S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford, 2013. URL: <https://books.google.com/books?id=koNqWR1uhPOC>.
- 2 Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In Martin Hirt and Adam Smith, editors, *Theory of Cryptography*, pages 635–658, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- 3 Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, page 1–10, New York, NY, USA, 2014. Association for Computing Machinery. doi:10.1145/2591796.2591877.
- 4 Yuval Dagan and Gil Kur. A bounded-noise mechanism for differential privacy, 2020. arXiv:2012.03817.
- 5 Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaudey, editor, *Advances in Cryptology - EUROCRYPT 2006*, pages 486–503, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

- 6 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, pages 265–284, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- 7 Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil Vadhan. On the complexity of differentially private data release: Efficient algorithms and hardness results. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC '09, page 381–390, New York, NY, USA, 2009. Association for Computing Machinery. doi:10.1145/1536414.1536467.
- 8 Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407, 2014. doi:10.1561/0400000042.
- 9 Badih Ghazi, Ravi Kumar, and Pasin Manurangsi. On avoiding the union bound when answering multiple differentially private queries, 2020. arXiv:2012.09116.
- 10 Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 61–70, October 2010. doi:10.1109/FOCS.2010.85.
- 11 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, page 705–714, New York, NY, USA, 2010. Association for Computing Machinery. doi:10.1145/1806689.1806786.
- 12 N.L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*. John Wiley & Sons Incorporated, 1995. URL: <https://books.google.com/books?id=q03oAAAACAAJ>.
- 13 D.A. Klain, G.C. Rota, and L.A.R. di Brozolo. *Introduction to Geometric Probability*. Lezioni Lincee. Cambridge University Press, 1997. URL: <https://books.google.com/books?id=Q1ytkNM6BtAC>.
- 14 Fang Liu. Generalized gaussian mechanism for differential privacy. *IEEE Transactions on Knowledge and Data Engineering*, 31:747–756, 2019.
- 15 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, page 94–103, USA, 2007. IEEE Computer Society. doi:10.1109/FOCS.2007.41.
- 16 Ilya Mironov. Rényi differential privacy. *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275, 2017.
- 17 Saralees Nadarajah. A generalized normal distribution. *Journal of Applied Statistics*, 32(7):685–694, 2005. doi:10.1080/02664760500079464.
- 18 Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2), 2017. doi:10.29012/jpc.v7i2.648.

A Deferred Proofs

A.1 Proof of Theorem 19

Proof of Theorem 19. The mechanism is given by modifying the NumericSparse algorithm given as Algorithm 3 in [8] by outputting 0 instead of \perp or 0 for all remaining queries instead of halting prematurely. The first two properties follow from the associated proofs in that text.

The third property follows because for all entries of \tilde{d} that SV does not output as 0 (for which the error, i.e. corresponding entry of $\tilde{d} - d$, is of course bounded by $\|d\|_\infty$), the error is drawn from $Lap(b)$ where $b = O(\sqrt{k \log(1/\delta_{SV})}/\epsilon_{SV})$. So the maximum error for these (at most $c_{SV} \leq k$) entries is stochastically dominated by the maximum of the absolute value of k of these Laplace random variables, which is at most tb with probability ke^{-t} . ◀

A.2 Proof of Theorem 1

We first need the following corollary of Lemma 15:

► **Corollary 21.** *Let x be chosen uniformly at random from a k -dimensional ℓ_p -sphere with arbitrary radius for $p \geq 1$. Then we have:*

$$\mathbb{E}[\|x\|_\infty] \leq \frac{5 \log^{1/p} k}{k^{1/p}} \|x\|_p$$

Proof. Since $\|x\|_\infty/\|x\|_p$ takes values in $[0, 1]$, by Lemma 15 we have:

$$\begin{aligned} \mathbb{E}[\|x\|_\infty/\|x\|_p] &= \int_0^1 \Pr[\|x\|_\infty/\|x\|_p \geq r] dr \\ &\leq \int_0^{\frac{2^{1+1/p} \log^{1/p} k}{k^{1/p}}} 1 dr + \int_{\frac{2^{1+1/p} \log^{1/p} k}{k^{1/p}}}^1 k \cdot \exp\left(-\frac{(k-1)r^p}{p}\right) dr \\ &\leq \frac{2^{1+1/p} \log^{1/p} k}{k^{1/p}} + \int_{\frac{2^{1+1/p} \log^{1/p} k}{k^{1/p}}}^1 k \cdot \exp\left(-\frac{(k-1)2^{p+1} \log k}{kp}\right) dr \\ &\leq \frac{2^{1+1/p} \log^{1/p} k}{k^{1/p}} + \int_{\frac{2^{1+1/p} \log^{1/p} k}{k^{1/p}}}^1 k \cdot \exp(-2 \log k) dr \\ &\leq \frac{2^{1+1/p} \log^{1/p} k}{k^{1/p}} + \frac{1}{k} \\ &\leq \frac{5 \log^{1/p} k}{k^{1/p}}. \end{aligned}$$

Here we use that $2^p \geq p$ for all $p \geq 1$ and that $(1 - \frac{c}{x})^x \leq e^{-c}$ for all $c \geq 0$. ◀

Proof of Theorem 1. We use Theorem 17 after rounding p up to the nearest even integer (this loses at most a multiplicative constant in the resulting error bounds). If the constant hidden in $\Theta(\log \log k)$ is a sufficiently large function of c_1 , this gives the desired tail bound, up to the additive $e^{-.001k/p}$ in the probability bound (which may be larger than the $e^{-t^p \log k}$ term for large values of p). To remove the additive $e^{-.001k/p}$: if the less than $e^{-.001k/p} \leq \delta$ probability event that $(\|x\|_p/\sigma)^p$ exceeds $2k/p$ occurs, we can instead just output $\tilde{d} = d$, i.e. instead set $x = 0$. This gives an $(\epsilon, 2\delta)$ -private mechanism that always satisfies $(\|x\|_p/\sigma)^p \leq 2k/p$, and then we can rescale our choice of δ appropriately. The tail bound can now be derived as in the proof of Theorem 17. Similarly, since we always have $(\|x\|_p/\sigma)^p \leq 2k/p$, the expectation of $\|x\|_\infty$ follows from Corollary 21. Finally, the expectation of $\|x\|_q$ for $1 \leq q \leq p$ follows by using Jensen's inequality twice and the unconditional upper bound on $\|x\|_p^p$:

$$\begin{aligned} \mathbb{E}[\|x\|_q] &\leq \mathbb{E}[\|x\|_q^{q/2}]^{2/q} = k^{1/q} \mathbb{E}[|x_1|^q]^{1/2} \leq k^{1/q} \mathbb{E}[|x_1|^p]^{1/2} = k^{1/q-1/2} \mathbb{E}[\|x\|_p^p] \\ &\leq k^{1/q-1/2} \cdot (2k/p)^{1/2} \sigma = O(k^{1/q} \sigma). \end{aligned}$$

A.3 Proof of Theorem 2

Proof of Theorem 2. The tail bound in Theorem 2 follows immediately from Theorem 20 by choosing p to be an even integer satisfying $p = \Theta(\log \log \log k)$.

For the expectation, we use the tail bound of Theorem 2. We have:

$$\begin{aligned}\mathbb{E}_{\tilde{d} \sim \mathcal{M}(d)} [|\tilde{d} - d|_\infty] &= \int_0^\infty \Pr[|\tilde{d} - d|_\infty \geq s] ds \\ &= \int_0^a \Pr[|\tilde{d} - d|_\infty \geq s] ds + \int_a^b \Pr[|\tilde{d} - d|_\infty \geq s] ds + \int_b^\infty \Pr[|\tilde{d} - d|_\infty \geq s] ds.\end{aligned}$$

We choose $a = \frac{2c\sqrt{k \log \log \log k \log(1/\delta)}}{\epsilon}$, $b = \frac{k\sqrt{\log(1/\delta)}}{\epsilon}$. The integral over $[0, a]$ is of course bounded by a . By Theorem 20, the integral over $[a, b]$ is bounded by $b \cdot e^{-\log^2 k} \leq \frac{\sqrt{\log(1/\delta)}}{\epsilon} \leq a$. Finally, to bound the third term, recall that the mechanism of Theorem 20 outputs d (i.e. effectively chooses $x, \hat{x} = 0$ instead) if $\|x\|_p$ is too large. So, unconditionally we have:

$$\|x\|_\infty \leq \|x\|_p \leq (2k/p)^{1/p} \sigma \leq \frac{2c\sqrt{k \log \log \log k \log(1/\delta)}}{\epsilon} \leq b.$$

So by the third property in Theorem 19 we have for $s \in [b, \infty)$:

$$\Pr_{\tilde{d} \sim \mathcal{M}(d)} [|\tilde{d} - d|_\infty \geq s] = \Pr_{x, \hat{x}} [\|x - \hat{x}\|_\infty \geq s] \leq ke^{-\Omega(s/(\sqrt{k \log(1/\delta)}/\epsilon))}.$$

And so by change of variables, with $s' = s/(\sqrt{k \log(1/\delta)}/\epsilon)$:

$$\int_b^\infty \Pr[|\tilde{d} - d|_\infty \geq s] ds \leq \frac{\sqrt{k \log(1/\delta)}}{\epsilon} \int_{\sqrt{k}}^\infty ke^{-\Omega(s')} ds' \leq \frac{k^{1.5} \sqrt{\log(1/\delta)}}{\epsilon} \cdot e^{-\Omega(\sqrt{k})} \leq a.$$

So we conclude

$$\mathbb{E}_{\tilde{d} \sim \mathcal{M}(d)} [|\tilde{d} - d|_\infty] \leq 3a = O\left(\frac{\sqrt{k \log \log \log k \log(1/\delta)}}{\epsilon}\right),$$

as desired. ◀

B Concentration of Generalized Gammas

In this section we consider the Generalized Gamma random variable $GGamma(a, b)$ parameterized by a, b with pdf:

$$p(x) = \frac{bx^{a-1}e^{-x/b}}{\Gamma(a/b)}, x \in (0, \infty).$$

Where the Gamma function $\Gamma(x)$ is defined over the positive reals as

$$\Gamma(z) = \int_0^\infty x^{z-1}e^{-x} dx.$$

We recall that $\Gamma(z)$ is a continuous analog of the factorial in that it satisfies $\Gamma(x+1) = x \cdot \Gamma(x)$. When $b = 1$, $GGamma(a, b)$ is exactly the Gamma random variable $Gamma(a)$ (we will use $Gamma$ to denote the random variable and Γ to denote the function to avoid ambiguous notation).

We want to show that sums of $GGamma(\frac{1}{p-1}, \frac{p}{p-1})$ random variables concentrate nicely. To do this, we will show that they are sub-gamma:

To show that $GGamma(\frac{1}{p-1}, \frac{p}{p-1})$ are sub-gamma, we will relate the moment-generating function of $GGamma(\frac{1}{p-1}, \frac{p}{p-1})$ to that of the Gamma random variable with the same mean using the following facts:

► **Fact 22.** For a Generalized Gamma random variable $X \sim G\text{Gamma}(a, b)$ the moments are $\mathbb{E}[X^r] = \frac{\Gamma((a+r)/b)}{\Gamma(a/b)}$. In particular, for a Gamma random variable $X \sim \text{Gamma}(a)$ the moments are $\mathbb{E}[X^r] = \frac{\Gamma(a+r)}{\Gamma(a)}$.

See e.g. Section 17.8.7 of [12] for a derivation of this fact. Note here that $G\text{Gamma}(\frac{1}{p-1}, \frac{p}{p-1})$ has mean $\mu = 1/\Gamma(1/p)$. To relate the moments of Generalized Gamma random variables to Gamma random variables' we note the following about μ :

► **Fact 23.** For all $p \geq 2$, we have $\frac{1}{p} \leq \frac{1}{\Gamma(1/p)} \leq \frac{1.2}{p}$.

Putting it all together, we get the following lemmas, which combined with Fact 23 give us Lemma 13:

► **Lemma 24.** Let $Y = G\text{Gamma}(\frac{1}{p-1}, \frac{p}{p-1})$ for $p \geq 2$. Then, for $\mu = \mathbb{E}[Y] = \frac{1}{\Gamma(1/p)}$, we have $Y \in \Gamma^+(\mu, 1)$.

Proof. We compare the moment-generating function of (the centered version of) Y to that of $X = \text{Gamma}(\mu)$ where $\mu = \mathbb{E}[Y]$. X is in $\Gamma(\mu, 1)$ so it suffices to show Y 's moment generating function is smaller than X 's. First, looking at the moment generating function of Y , we have:

$$\begin{aligned} \mathbb{E}[e^{\lambda Y}] &= 1 + \lambda\mu + \sum_{r=2}^{\infty} \left[\frac{\lambda^r}{r!} \mathbb{E}[Y^r] \right] \\ &= 1 + \lambda\mu + \sum_{r=2}^{\infty} \left[\frac{\lambda^r \Gamma(\frac{1}{p} + \frac{r(p-1)}{p})}{r! \Gamma(\frac{1}{p})} \right] \\ &\stackrel{(a)}{\leq} 1 + \lambda\mu + \sum_{r=2}^{\infty} \left[\frac{\lambda^r \Gamma(\frac{1}{p} + r)}{r! \Gamma(\frac{1}{p})} \right] \\ &\stackrel{(b)}{\leq} 1 + \lambda\mu + \sum_{r=2}^{\infty} \left[\frac{\lambda^r \Gamma(\mu + r)}{r! \Gamma(\mu)} \right] = \mathbb{E}[e^{\lambda X}]. \end{aligned}$$

(a) follows because the Gamma function is monotonically increasing in the range $[1.5, \infty)$. (b) follows because $\mu = \frac{1}{\Gamma(1/p)} \geq 1/p$ for $p \geq 1$, and because for positive integers r , $\frac{\Gamma(x+r)}{\Gamma(x)} = \prod_{i=0}^{r-1} (x+i)$ is monotonically increasing in x . Since $X \in \Gamma^+(\mu, 1)$ and X, Y have the same mean, we have that $Y \in \Gamma^+(\mu, 1)$ as well. ◀

► **Lemma 25.** Let $Y = G\text{Gamma}(\frac{1}{p-1}, \frac{p}{p-1})$ for $p \geq 3$. Then, for $\mu = \mathbb{E}[Y] = \frac{1}{\Gamma(1/p)}$, we have $Y \in \Gamma^-(\mu, 3/2)$.

Proof. Similarly to the previous lemma, we have for all $0 \leq \lambda \leq 2/3$:

$$\begin{aligned} \mathbb{E}[e^{-\lambda Y}] &= 1 - \lambda\mu + \sum_{r=2}^{\infty} \left[\frac{(-\lambda)^r \Gamma(\frac{1}{p} + \frac{r(p-1)}{p})}{r! \Gamma(\frac{1}{p})} \right] \\ &= 1 - \lambda\mu + \sum_{r=1}^{\infty} \left[\frac{\lambda^{2r}}{(2r)!} \cdot \frac{\Gamma(\frac{1}{p} + 2r \frac{p-1}{p})}{\Gamma(\frac{1}{p})} \left(1 - \frac{\lambda}{2r+1} \cdot \frac{\Gamma(\frac{1}{p} + (2r+1) \frac{p-1}{p})}{\Gamma(\frac{1}{p} + 2r \frac{p-1}{p})} \right) \right] \\ &= 1 - \lambda\mu + \sum_{r=1}^{\infty} \left[\frac{\lambda^{2r}}{(2r)!} \cdot \frac{\Gamma(\frac{1}{p} + 2r)}{\Gamma(\frac{1}{p})} \left(\frac{\Gamma(\frac{1}{p} + 2r \frac{p-1}{p})}{\Gamma(\frac{1}{p} + 2r)} - \frac{\lambda}{2r+1} \cdot \frac{\Gamma(\frac{1}{p} + (2r+1) \frac{p-1}{p})}{\Gamma(\frac{1}{p} + 2r)} \right) \right] \end{aligned}$$

$$\begin{aligned}
 \dots &\stackrel{(c)}{\leq} 1 - \lambda\mu + \sum_{r=1}^{\infty} \left[\frac{\lambda^{2r}}{(2r)!} \cdot \frac{\Gamma(\frac{1}{p} + 2r)}{\Gamma(\frac{1}{p})} \left(1 - \frac{\lambda}{2r+1} \cdot \frac{\Gamma(\frac{1}{p} + 2r + 1)}{\Gamma(\frac{1}{p} + 2r)} \right) \right] \\
 &\stackrel{(d)}{\leq} 1 - \lambda\mu + \sum_{r=1}^{\infty} \left[\frac{\lambda^{2r}}{(2r)!} \cdot \frac{\Gamma(\mu + 2r)}{\Gamma(\mu)} \left(1 - \frac{\lambda}{2r+1} \cdot \frac{\Gamma(\mu + 2r + 1)}{\Gamma(\mu + 2r)} \right) \right] \\
 &= 1 - \lambda\mu + \sum_{r=2}^{\infty} \left[\frac{(-\lambda)^r}{r!} \cdot \frac{\Gamma(\mu + r)}{\Gamma(\mu)} \right] = \mathbb{E}[e^{-\lambda X}].
 \end{aligned}$$

Which, up to proving (c), (d) hold, shows that $Y \in \Gamma^-(\mu, 3/2)$ since X and Y have the same mean and $X \in \Gamma^-(\mu, 0) \subset \Gamma^-(\mu, 3/2)$. (c) follows because the change in each term in the sum is

$$\begin{aligned}
 &\frac{\lambda^{2r}}{(2r)!} \frac{1}{\Gamma(\frac{1}{p})} \\
 &\left[\Gamma\left(\frac{1}{p} + 2r\right) - \Gamma\left(\frac{1}{p} + 2r \frac{p-1}{p}\right) - \frac{\lambda}{2r+1} \left(\Gamma\left(\frac{1}{p} + 2r + 1\right) - \Gamma\left(\frac{1}{p} + (2r+1) \frac{p-1}{p}\right) \right) \right].
 \end{aligned}$$

To show this expression is non-negative, it suffices to show that just the term in the brackets is positive, or equivalently, for all $r \geq 2, p \geq 3$:

$$\Gamma\left(\frac{1}{p} + 2r\right) \left(1 - \frac{\Gamma\left(\frac{1}{p} + 2r \frac{p-1}{p}\right)}{\Gamma\left(\frac{1}{p} + 2r\right)} \right) \geq \frac{\lambda}{2r+1} \Gamma\left(\frac{1}{p} + 2r + 1\right) \left(1 - \frac{\Gamma\left(\frac{1}{p} + (2r+1) \frac{p-1}{p}\right)}{\Gamma\left(\frac{1}{p} + 2r + 1\right)} \right).$$

Since we have $\Gamma\left(\frac{1}{p} + 2r + 1\right) = \left(\frac{1}{p} + 2r\right)\Gamma\left(\frac{1}{p} + 2r\right) \leq (2r+1)\left(\frac{1}{p} + 2r\right)$, it further suffices to just show:

$$f(r, p) := \frac{\left(1 - \frac{\Gamma\left(\frac{1}{p} + 2r \frac{p-1}{p}\right)}{\Gamma\left(\frac{1}{p} + 2r\right)} \right)}{\left(1 - \frac{\Gamma\left(\frac{1}{p} + (2r+1) \frac{p-1}{p}\right)}{\Gamma\left(\frac{1}{p} + 2r + 1\right)} \right)} \geq \lambda.$$

For any fixed $r \geq 2$, one can verify analytically that $f(r, p)$ is monotonically decreasing in p over $p \in [1, \infty)$ and the limit as p goes to infinity is $g(r) := \frac{2r\psi(2r)}{(2r+1)\psi(2r+1)}$ where ψ is the digamma function $\psi(x) = \frac{d}{dx} \Gamma(x)$. One can also verify analytically that $g(r)$ is monotonically increasing, and $g(2) \approx .6672$. So, for all $r \geq 2, p \geq 3$ we have $f(r, p) > 2/3$ and thus for $\lambda \in [0, 2/3]$, the inequality (c) is satisfied.

(d) follows by looking at the function

$$z(x) = \frac{\Gamma(x+r)}{\Gamma(x)} \left(1 - \frac{\lambda}{r+1} \cdot \frac{\Gamma(x+r+1)}{\Gamma(x+r)} \right) = \left(1 - \frac{\lambda(x+r)}{r+1} \right) \prod_{i=0}^{r-1} (x+i).$$

For $r \geq 2, \lambda \leq 1$, one can verify analytically that $z(x)$ is monotonically increasing in the interval $(0, 1/2] \supseteq (0, \frac{1+2}{p}] \supseteq (0, \mu]$. Since $\mu \geq \frac{1}{p}$, this gives that each term in the right-hand-side of (d) is larger than the corresponding term on the left-hand-side. ◀