

Differential Privacy on Finite Computers*

Victor Balcer[†] Salil Vadhan[‡]

Center for Research on Computation & Society

School of Engineering & Applied Sciences

Harvard University

vbalcer@g.harvard.edu, salil_vadhan@harvard.edu

December 24, 2018

Abstract

We consider the problem of designing and analyzing differentially private algorithms that can be implemented on *discrete* models of computation in *strict* polynomial time, motivated by known attacks on floating point implementations of real-arithmetic differentially private algorithms (Mironov, CCS 2012) and the potential for timing attacks on expected polynomial-time algorithms. As a case study, we examine the basic problem of approximating the histogram of a categorical dataset over a possibly large data universe \mathcal{X} . The classic Laplace Mechanism (Dwork, McSherry, Nissim, Smith, TCC 2006 and J. Privacy & Confidentiality 2017) does not satisfy our requirements, as it is based on real arithmetic, and natural discrete analogues, such as the Geometric Mechanism (Ghosh, Roughgarden, Sundarajan, STOC 2009 and SICOMP 2012), take time at least linear in $|\mathcal{X}|$, which can be exponential in the bit length of the input.

In this paper, we provide strict polynomial-time discrete algorithms for approximate histograms whose simultaneous accuracy (the maximum error over all bins) matches that of the Laplace Mechanism up to constant factors, while retaining the same (pure) differential privacy guarantee. One of our algorithms produces a sparse histogram as output. Its “per-bin accuracy” (the error on individual bins) is worse than that of the Laplace Mechanism by a factor of $\log |\mathcal{X}|$, but we prove a lower bound showing that this is necessary for any algorithm that produces a sparse histogram. A second algorithm avoids this lower bound, and matches the per-bin accuracy of the Laplace Mechanism, by producing a compact and efficiently computable representation of a dense histogram; it is based on an $(n + 1)$ -wise independent implementation of an appropriately clamped version of the Discrete Geometric Mechanism.

1 Introduction

Differential Privacy [DMNS06] is by now a well-established framework for privacy-protective statistical analysis of sensitive datasets. Much work on differential privacy involves an interplay between statistics and computer science. Statistics provides many of the (non-private) analyses that we wish to approximate with differentially private algorithms, as well as probabilistic tools that are useful in analyzing such algorithms, which are necessarily randomized. From computer science, differential privacy draws upon a tradition of adversarial modeling and strong security definitions,

*A condensed version of this paper appeared in ITCS 2018 [BV18]

[†]Supported by NSF grant CNS-1237235 and CNS-1565387.

[‡]salil.seas.harvard.edu. Supported by NSF grant CNS-1237235, a Simons Investigator Award, and a grant from the Sloan Foundation.

techniques for designing and analyzing randomized algorithms, and considerations of algorithmic resource constraints (such as time and memory).

Because of its connection to statistics, it is very natural that much of the literature on differential privacy considers the estimation of real-valued functions on real-valued data (e.g. the sample mean) and introduces noise from continuous probability distributions (e.g. the Laplace distribution) to obtain privacy. However, these choices are incompatible with standard computer science models for algorithms (like the Turing machine or RAM model) as well as implementation on physical computers (which use only finite approximations to real arithmetic, e.g. via floating point numbers). This discrepancy is not just a theoretical concern; Mironov [Mir12] strikingly demonstrated that common floating-point implementations of the most basic differentially private algorithm (the Laplace Mechanism) are vulnerable to real attacks. Mironov shows how to prevent his attack with a simple modification to the implementation, but this solution is specific to a single differentially private mechanism and particular floating-point arithmetic standard. His solution increases the error by a constant factor and seems likely to be quite efficient in practice. However, he provides no bounds on asymptotic running time. Gazeau, Miller and Palamidessi [GMP13] provide more general conditions under which an implementation of real numbers and a mechanism that perturbs the correct answer with noise maintains differential privacy. However, they do not provide an explicit construction with bounds on accuracy and running time.

From a theoretical point of view, a more appealing approach to resolving these issues is to avoid real or floating-point arithmetic entirely and only consider differentially private computations that involve discrete inputs and outputs, and rational probabilities, as first done in [DKM⁺06]. Such algorithms are realizable in standard discrete models of computation. However, some such algorithms have running times that are only bounded in expectation (e.g. due to sampling from an exponential distribution supported on the natural numbers), and this raises a potential vulnerability to timing attacks. If an adversary can observe the running time of the algorithm, it learns something about the algorithm’s coin tosses, which are assumed to be secret in the definition of differential privacy. (Even if the time cannot be directly observed, in practice an adversary can determine an upper bound on the running time, which again is information that is implicitly assumed to be secret in the privacy definition.)

Because of these considerations, we advocate the following principle:

Differential Privacy for Finite Computers:

*We should describe how to implement differentially private algorithms on **discrete** models of computation with **strict** bounds on running time (ideally polynomial in the bit length of the input) and **analyze** the effects of those constraints on both privacy and accuracy.*

Note that a strict *bound* on running time does not in itself prevent timing attacks, but once we have such a bound, we can pad all executions to take the same amount of time. Also, while standard discrete models of computation (e.g. randomized Turing machines) are defined in terms of countable rather than finite resources (e.g. the infinite tape), if we have a strict bound on running time, then once we fix an upper bound on input length, they can indeed be implemented on a truly finite computer (e.g. like a randomized Boolean circuit).

In many cases, the above goal can be achieved by appropriate discretizations and truncations applied to a standard, real-arithmetic differentially private algorithm. However, such modifications can have a nontrivial price in accuracy or privacy, and thus we also call for a rigorous analysis of these effects.

In this paper, we carry out a case study of achieving “differential privacy for finite computers” for one of the first tasks studied in differential privacy, namely approximating a histogram of a

categorical dataset. Even this basic problem turns out to require some nontrivial effort, particularly to maintain strict polynomial time, optimal accuracy and pure differential privacy when the data universe is large.

We recall the definition of differential privacy.

Definition 1.1 ([DMNS06]). Let $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{R}$ be a randomized algorithm. We say \mathcal{M} is (ε, δ) -**differentially private** if for every pair of datasets D and D' that differ on one row and every subset $S \subseteq \mathcal{R}$

$$\Pr[\mathcal{M}(D) \in S] \leq e^\varepsilon \cdot \Pr[\mathcal{M}(D') \in S] + \delta$$

We say an (ε, δ) -differentially private algorithm satisfies **pure differential privacy** when $\delta = 0$ and say it satisfies **approximate differential privacy** when $\delta > 0$.

In this paper, we study the problem of estimating the *histogram* of a dataset $D \in \mathcal{X}^n$, which is the vector $c = c(D) \in \mathbb{N}^{\mathcal{X}}$, where c_x is the number of rows in D that have value x . Histograms can be approximated while satisfying differential privacy using the *Laplace Mechanism*, introduced in the original paper of Dwork, McSherry, Nissim and Smith [DMNS06]. Specifically, to obtain $(\varepsilon, 0)$ -differential privacy, we can add independent noise distributed according to a Laplace distribution, specifically $\text{Lap}(2/\varepsilon)$, to each component of c and output the resulting vector \tilde{c} . Here $\text{Lap}(2/\varepsilon)$ is the *continuous*, real-valued random variable with probability density function $f(z)$ that is proportional to $\exp(-\varepsilon \cdot |z|/2)$. The Laplace Mechanism also achieves very high accuracy in two respects:

Per-Query Error: For each bin $x \in \mathcal{X}$, with high probability we have $|\tilde{c}_x - c_x| \leq O(1/\varepsilon)$.

Simultaneous Error: With high probability, we have $\max_x |\tilde{c}_x - c_x| \leq O(\log(|\mathcal{X}|)/\varepsilon)$.

Note that both of the bounds are independent of the number n of rows in the dataset, and so the fractional error vanishes linearly as n grows.

Simultaneous error is the more well-studied notion in the differential privacy literature, but we consider per-query error to be an equally natural concept: if we think of the approximate histogram \tilde{c} as containing approximate answers to the $|\mathcal{X}|$ different counting queries corresponding to the bins of \mathcal{X} , then per-query error captures the error as experienced by an analyst who may be only interested in one or a few of the bins of \tilde{c} . The advantage of considering per-query error is that it can be significantly smaller than the simultaneous error, as is the case in the Laplace Mechanism when the data universe \mathcal{X} is very large. It is known that both of the error bounds achieved by the Laplace Mechanism are optimal up to constant factors; no $(\varepsilon, 0)$ -differentially private algorithm for histograms can achieve smaller per-query error or simultaneous error [HT10, BBKN14].

Unfortunately, the Laplace Mechanism uses real arithmetic and thus cannot be implemented on a finite computer. To avoid real arithmetic, we could use the Geometric Mechanism [GRS12], which adds noise to each component of c according to the 2-sided geometric distribution, $\text{Geo}(2/\varepsilon)$, which is supported on the integers and has probability mass function $f(z) \propto \exp(-\varepsilon \cdot |z|/2)$. However, this mechanism uses integers of unbounded size and thus cannot be implemented on a finite computer. Indeed, while the algorithm can be implemented with a running time that is bounded in expectation (after reducing ε so that $e^{\varepsilon/2}$ and hence all the probabilities are rational numbers), truncating long executions or allowing an adversary to observe the actual running time can lead to a violation of differential privacy. Thus, as first described by Dwork, Kenthapadi, McSherry, Mironov and Naor [DKM⁺06], it is better to restrict the output of the mechanism to a binary representation of fixed length in order to avoid small tail probabilities. Similarly, we work with the Truncated Geometric Mechanism of Ghosh, Roughgarden and Sundararajan [GRS12], where we clamp each noisy count

\tilde{c}_x to the interval $[0, n]$. We observe that the resulting probability distribution of \tilde{c}_x , supported on $\{0, 1, \dots, n\}$, can be described explicitly in terms of c_x , ε and n , and it can be sampled in polynomial time using only integer arithmetic (after ensuring $e^{\varepsilon/2}$ is rational). Thus, we obtain:

Theorem 1.2 (Bounded Geometric Mechanism, informal statement of Thm. 5.3). *For every finite \mathcal{X} , n and $\varepsilon \in (0, 1]$, there is an $(\varepsilon, 0)$ -differentially private algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \{0, 1, \dots, n\}^{\mathcal{X}}$ for histograms achieving:*

- Per-query error $O(1/\varepsilon)$.
- Simultaneous error $O(\log(|\mathcal{X}|)/\varepsilon)$.
- Strict running time $\tilde{O}(|\mathcal{X}|/\varepsilon) \cdot \log^2 n + O(n \log n \cdot \log |\mathcal{X}|)$.

We note that while we only consider our particular definition of per-query accuracy, namely that with high probability $|\tilde{c}_x - c_x| \leq O(1/\varepsilon)$, Ghosh, Roughgarden and Sundararajan [GRS12] proved that the output of the Bounded Geometric Mechanism can be used (with post-processing) to get optimal expected loss with respect to an extremely general class of loss functions and arbitrary priors. The same result applies to each individual noisy count \tilde{c}_x output by our mechanism, since each bin is distributed according to the Bounded Geometric Mechanism (up to a modification of ε to ensure rational probabilities).

The Bounded Geometric Mechanism is not polynomial time for large data universes \mathcal{X} . Indeed, its running time (and output length) is linear in $|\mathcal{X}|$, rather than polynomial in the bit length of data elements, which is $\log |\mathcal{X}|$. To achieve truly polynomial time, we can similarly discretize and truncate a variant of the Stability-Based Histogram that was introduced by Korolova, Kenthapadi, Mishra and Ntoulas [KKMN09], and explicitly described by Bun, Nissim and Stemmer [BNS16]. This mechanism only adds $\text{Lap}(2/\varepsilon)$ noise to the *nonzero* components of c_x and then retains only the noisy values \tilde{c}_x that are larger than a threshold $t = \Theta(\log(1/\delta)/\varepsilon)$. Thus, the algorithm only outputs a partial histogram, i.e. counts \tilde{c}_x for a subset of the bins x , with the rest of the counts being treated as zero. By replacing the use of the Laplace Mechanism with the (rational) Bounded Geometric Mechanism as above, we can implement this algorithm in strict polynomial time:

Theorem 1.3 (Stability-Based Histogram, informal statement of Thm. 5.6). *For every finite \mathcal{X} , n , $\varepsilon \in (0, 1]$ and $\delta \in (0, 1/n)$, there is an (ε, δ) -differentially private algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \{0, 1, \dots, n\}^{\subseteq \mathcal{X}}$ for histograms achieving:*

- Per-query error $O(1/\varepsilon)$ on bins with true count at least $O(\log(1/\delta)/\varepsilon)$.
- Simultaneous error $O(\log(1/\delta)/\varepsilon)$.
- Strict running time $\tilde{O}((n/\varepsilon) \cdot \log(1/\delta)) + O(n \log n \cdot \log |\mathcal{X}|)$.

Notice that the simultaneous error bound of $O(\log(1/\delta)/\varepsilon)$ is better than what is achieved by the Laplace Mechanism when $\delta > 1/|\mathcal{X}|$, and is known to be optimal up to constant factors in this range of parameters (see Theorem 7.1). The fact that this error bound is independent of the data universe size $|\mathcal{X}|$ makes it tempting to apply even for infinite data domains \mathcal{X} . However, we note that when \mathcal{X} is infinite, it is impossible for the algorithm to have a strict bound on running time (as it needs time to read arbitrarily long data elements) and thus is vulnerable to timing attacks and is not implementable on a finite computer. Note also that the per-query error bound only holds on bins with large enough true count (namely, those larger than our threshold t); we will discuss this point further below.

A disadvantage of the Stability-based Histogram is that it sacrifices pure differential privacy. It is natural to ask whether we can achieve polynomial running time while retaining pure differential privacy. A step in this direction was made by Cormode, Procopiuc, Srivastava and Tran [CPST12]. They observe that for an appropriate threshold $t = \Theta(\log(|\mathcal{X}|)/\varepsilon)$, if we run the Bounded Geometric Mechanism and only retain the noisy counts \tilde{c}_x that are larger than t , then the expected number of bins that remain is less than $n+1$. Indeed, the expected number of bins we retain whose true count is zero (“empty bins”) is less than 1. They describe a method to directly sample the distribution of the empty bins that are retained, without actually adding noise to all $|\mathcal{X}|$ bins. This yields an algorithm whose output length is polynomial in expectation. However, the output length is not strictly polynomial, as there is a nonzero probability of outputting all $|\mathcal{X}|$ bins. And it is not clear how to implement the algorithm even in expected polynomial time, because even after making the probabilities rational, they have denominators of bit length linear in $|\mathcal{X}|$.

To address these issues, we consider a slightly different algorithm. Instead of trying to retain all noisy counts \tilde{c}_x that are larger than some fixed threshold t , we retain the n largest noisy counts (since there are at most n nonzero true counts). This results in a mechanism whose output length is always polynomial, rather than only in expectation. However, the probabilities still have denominators of bit length linear in $|\mathcal{X}|$. Thus, we show how to approximately sample from this distribution, to within an arbitrarily small statistical distance δ , at the price of a $\text{poly}(\log(1/\delta))$ increase in running time. Naively, this would result only in $(\varepsilon, O(\delta))$ -differential privacy. However, when δ is significantly smaller than $1/|\mathcal{R}|$, where \mathcal{R} is the range of the mechanism, we can convert an (ε, δ) -differentially private mechanism to an $(\varepsilon, 0)$ -differentially private mechanism by simply outputting a uniformly random element of \mathcal{R} with small probability. (A similar idea for the case that $|\mathcal{R}| = 2$ has been used in [KLN⁺11, CDK17].) Since our range is of at most exponential size (indeed at most polynomial in bit length), the cost in our runtime for taking $\delta \ll 1/|\mathcal{R}|$ is at most polynomial. With these ideas we obtain:

Theorem 1.4 (Pure DP Histogram in Polynomial Time, informal statement of Thm. 6.11). *For every finite \mathcal{X} , n and $\varepsilon \in (0, 1]$, there is an $(\varepsilon, 0)$ -differentially private algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \{0, 1, \dots, n\}^{\subseteq \mathcal{X}}$ for histograms achieving:*

- *Per-query error $O(1/\varepsilon)$ on bins with true count at least $O(\log(|\mathcal{X}|)/\varepsilon)$.*
- *Simultaneous error $O(\log(|\mathcal{X}|)/\varepsilon)$.*
- *Strict running time $\tilde{O}(n^2 \cdot \log^2 |\mathcal{X}| + n^2 \cdot \log(1/\varepsilon) + n \cdot \log |\mathcal{X}| \cdot \log(1/\varepsilon))$.*

It is an open problem as to whether or not one can improve the nearly quadratic dependence in running time on n to nearly linear while maintaining the sparsity, privacy and accuracy guarantees achieved in Theorem 1.4.

Both Theorems 1.3 and 1.4 only retain per-query error $O(1/\varepsilon)$ on bins with a large enough true count. We also prove a lower bound showing that this limitation is inherent in any algorithm that outputs a sparse histogram (as both of these algorithms do).

Theorem 1.5 (Lower Bound on Per-Query Error for Sparse Histograms, Theorem 7.2). *Suppose that there is an (ε, δ) -differentially private algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \{0, 1, \dots, n\}^{\mathcal{X}}$ for histograms that always outputs histograms with at most n' nonempty bins and has per-query error at most E on all bins. Then*

$$E \geq \Omega\left(\frac{\min\{\log |\mathcal{X}|, \log(1/\delta)\}}{\varepsilon}\right)$$

provided that $\varepsilon > 0$, $\varepsilon^2 > \delta > 0$ and $|\mathcal{X}| \geq (n')^2$.

This lower bound is similar in spirit to a lower bound of [BBKN14], which shows that no $(\varepsilon, 0)$ -differentially private PAC learner for “point functions” (functions that are 1 on exactly one element of the domain) can produce sparse functions as hypotheses.

To bypass this lower bound, we can consider algorithms that produce succinct descriptions of dense histograms. That is, the algorithm can output a polynomial-length description of a function $\tilde{c} : \mathcal{X} \rightarrow [0, n]$ that can be evaluated in polynomial time, even though \mathcal{X} may be of exponential size. We show that this relaxation allows us to regain per-query error $O(1/\varepsilon)$.

Theorem 1.6 (Polynomial-Time DP Histograms with Optimal Per-Query Accuracy, informal statement of Thm. 8.6). *For every finite \mathcal{X} , n and $\varepsilon \in (0, 1]$, there is an $(\varepsilon, 0)$ -differentially private algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{H}$ for histograms (where \mathcal{H} is an appropriate class of succinct descriptions of histograms) achieving:*

- Per-query error $O(1/\varepsilon)$.
- Simultaneous error $O(\log(|\mathcal{X}|)/\varepsilon)$.
- Strict running time $\tilde{O}((n/\varepsilon) \cdot \log |\mathcal{X}|)$.
- Evaluating a count takes time $\tilde{O}((n/\varepsilon) \cdot \log |\mathcal{X}|)$.¹

The algorithm is essentially an $(n+1)$ -wise independent instantiation of the Bounded Geometric Mechanism. Specifically, we release a function $h : \mathcal{X} \rightarrow \{0, 1\}^r$ selected from an $(n+1)$ -wise independent family of hash functions, and for each $x \in \mathcal{X}$, we view $h(x)$ as coin tosses specifying a sample from the Bounded Geometric Distribution. That is, we let $S : \{0, 1\}^r \rightarrow [0, n]$ be an efficient sampling algorithm for the Bounded Geometric Distribution, and then $\tilde{c}_x = S(h(x))$ is our noisy count for x . The hash function is chosen randomly from the family conditioned on values \tilde{c}_x for the nonempty bins x , which we obtain by running the actual Bounded Geometric Mechanism on those bins. The $(n+1)$ -wise independence ensures that the behavior on any two neighboring datasets (which together involve at most $n+1$ distinct elements of \mathcal{X}) are indistinguishable in the same way as in the ordinary Bounded Geometric Mechanism. The per-query accuracy comes from the fact that the marginal distributions of each of the noisy counts are the same as in the Bounded Geometric Mechanism.²

As far as we know, the only other use of limited independence in constructing differentially private algorithms is a use of pairwise independence by [BBKN14] in differentially private PAC learning algorithms for the class of point functions. Although that problem is related to the one we consider (releasing a histogram amounts to doing “query release” for the class of point functions, as discussed below), the design and analysis of our algorithm appears quite different. (In particular, our analysis seems to rely on $(n+1)$ -wise independence in an essential way.)

Another potential interest in our technique is as another method for bypassing limitations of *synthetic data* for *query release*. Here, we have a large family of predicates $\mathcal{Q} = \{q : \mathcal{X} \rightarrow \{0, 1\}\}$, and are interested in differentially private algorithms that, given a dataset $D = (x_1, \dots, x_n) \in \mathcal{X}^n$, output a “summary” $\mathcal{M}(D)$ that allows one to approximate the answers to all of the *counting queries* $q(D) = \sum_i q(x_i)$ associated with predicates $q \in \mathcal{Q}$. For example, if \mathcal{Q} is the family of *point functions* consisting of all predicates that evaluate to 1 on exactly one point in the data universe \mathcal{X} , then this query release problem amounts to approximating the histogram of D . The fundamental

¹In the original version of our paper [BV18], both the running time and evaluation time were missing logarithmic factors.

²Actually, we incur a small approximation error in matching the domain of the sampling procedure to the range of a family of hash functions.

result of Blum, Ligett, and Roth [BLR13] and successors show that this is possible even for families \mathcal{Q} and data universes \mathcal{X} that are of size exponential in n . Moreover, the summaries produced by these algorithms has the form of a synthetic dataset — a dataset $\hat{D} \in \mathcal{X}^{\hat{n}}$ such that for every query $q \in \mathcal{Q}$, we have $q(\hat{D}) \approx q(D)$. Unfortunately, it was shown in [UV11] that even for very simple families \mathcal{Q} of queries, such correlations between pairs of binary attributes, constructing such a differentially private synthetic dataset requires time exponential in the bit length $\log |\mathcal{X}|$ of data universe elements. Thus, it is important to find other ways of representing approximate answers to natural families \mathcal{Q} of counting queries, which can bypass the inherent limitations of synthetic data, and progress along these lines was made in a variety of works [GRU12, CKKL12, HRS12, TUV12, CTUW14, DNT15]. Our algorithm, and its use of $(n + 1)$ -wise independence, can be seen as yet another representation that bypasses a limitation of synthetic data (albeit a statistical rather than computational one). Indeed, a sparse histogram is simply a synthetic dataset that approximates answers to all point functions, and by Theorem 1.5, our algorithm achieves provably better per-query accuracy than is possible with synthetic datasets. This raises the question of whether similar ideas can also be useful in bypassing the computational limitations of synthetic data for more complex families of counting queries.

2 Preliminaries

Throughout this paper, let \mathbb{N} be the set $\{0, 1, \dots\}$, \mathbb{N}_+ be the set $\{1, 2, \dots\}$ and \mathbb{N}^{-1} be the set $\{1/n : n \in \mathbb{N}_+\}$. For $n \in \mathbb{N}_+$, let $[n]$ denote the set $\{0, \dots, n\}$ and $[n]_+$ denote the set $\{1, \dots, n\}$. (Notice that $|[n]| = n + 1$ while $|[n]_+| = n$.) Given a set A and finite set B , we define A^B to be the set of length $|B|$ vectors over A indexed by the elements of B .

2.1 Differential Privacy

We define a **dataset** $D \in \mathcal{X}^n$ to be an ordered tuple of $n \geq 1$ rows where each row is drawn from a discrete **data universe** \mathcal{X} with each row corresponding to an individual. Two datasets $D, D' \in \mathcal{X}^n$ are considered **neighbors** if they differ in exactly one row.

Definition 2.1 ([DMNS06]). Let $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{R}$ be a randomized algorithm. We say \mathcal{M} is (ε, δ) -**differentially private** if for every pair of neighboring datasets D and D' and every subset $S \subseteq \mathcal{R}$

$$\Pr[\mathcal{M}(D) \in S] \leq e^\varepsilon \cdot \Pr[\mathcal{M}(D') \in S] + \delta$$

We say an (ε, δ) -differentially private algorithm satisfies **pure differential privacy** when $\delta = 0$ and say it satisfies **approximate differential privacy** when $\delta > 0$. Intuitively, the ε captures an upper bound on an adversary’s ability to determine whether a particular individual is in the dataset. And the δ parameter represents an upper bound of the probability of a catastrophic privacy breach (e.g. the entire dataset is released). The common setting of parameters takes $\varepsilon \in (0, 1]$ to be a small constant and δ to be negligible in n .

The following properties of differentially private algorithms will be used in some of our proofs.

Lemma 2.2 (post-processing [DMNS06]). *Let $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ be (ε, δ) -differentially private and $T : \mathcal{Y} \rightarrow \mathcal{Z}$ be any randomized function. Then $T \circ \mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Z}$ is (ε, δ) -differentially private.*

Lemma 2.3 (group privacy [DMNS06]). *Let $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}$ be (ε, δ) -differentially private. Let $D_1, D_2 \subseteq \mathcal{X}^n$ be datasets such that D_2 can be obtained by changing at most m rows of D_1 . Then for all $S \subseteq \mathcal{Y}$*

$$\Pr[\mathcal{M}(D_1) \in S] \leq e^{m\varepsilon} \cdot \Pr[\mathcal{M}(D_2) \in S] + e^{m\varepsilon} \cdot \delta/\varepsilon$$

Lemma 2.4 (composition [DL09]). *Let $\mathcal{M}_1 : \mathcal{X}^n \rightarrow \mathcal{Y}_1$ be $(\varepsilon_1, \delta_1)$ -differentially private and $\mathcal{M}_2 : \mathcal{X}^n \rightarrow \mathcal{Y}_2$ be $(\varepsilon_2, \delta_2)$ -differentially private. Define $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{Y}_1 \times \mathcal{Y}_2$ as $\mathcal{M}(x) = (\mathcal{M}_1(x), \mathcal{M}_2(x))$. Then \mathcal{M} is $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$ -differentially private.*

2.2 Histograms

For $x \in \mathcal{X}$, the **point function** $c_x : \mathcal{X}^n \rightarrow \mathbb{N}$ is defined to count the number of occurrences of x in a given dataset, i.e.

$$c_x(D) = |\{i \in [n]_+ : D_i = x\}|$$

In this paper we focus on algorithms for privately releasing approximations to the values of all point functions, also known as a **histogram**. A histogram is a collection of **bins**, one for each element x in the data universe, with the x^{th} bin consisting of its **label** x and a **count** $c_x \in \mathbb{N}$.

2.2.1 Representations

The input to our algorithms is always a dataset (i.e. an element $D \in \mathcal{X}^n$) and the outputs represent approximate histograms. We consider the following histogram representations as our algorithms' outputs:

- A vector in $\mathbb{N}^{\mathcal{X}}$. We use $\{\tilde{c}_x\}_{x \in \mathcal{X}}$ to denote a histogram where $\tilde{c}_x \in \mathbb{N}$ is the approximate count for the element x .
- A partial vector $h \in (\mathcal{X} \times \mathbb{N})^*$ such that each element $x \in \mathcal{X}$ appears at most once in h with each pair $(x, \tilde{c}_x) \in \mathcal{X} \times \mathbb{N}$ interpreted as element x having approximate count \tilde{c}_x . Elements x not listed in the partial vector are assumed to have count $\tilde{c}_x = 0$. Implicitly, an algorithm can return a partial vector by releasing bins for a subset of \mathcal{X} .³
- A data structure, encoded as a string, which defines a function $h : \mathcal{X} \rightarrow \mathbb{N}$ where $h(x)$, denoted h_x , is the approximate count for $x \in \mathcal{X}$ and h_x is efficiently computable given this data structure (e.g. time polynomial in the length of the data structure). In Section 8, this data structure consists of the coefficients of a polynomial, along with some parameters.

Each representation is able to express any histogram over \mathcal{X} . The difference between them is the memory used and the efficiency of computing a count. For example, computing the approximate count for $x \in \mathcal{X}$, when using the data structure representation is bounded by the time it takes to compute the associated function. But when using partial vectors, one only needs to iterate through the vector to determine the approximate count.

We define the following class of histograms. Let $\mathcal{H}_{n,n'}(\mathcal{X}) \subseteq \mathbb{N}^{\mathcal{X}}$ be the set of all histograms over \mathcal{X} with integer counts in $[0, n]$ (or \mathbb{N} when $n = \infty$) and at most n' of them nonzero. By using partial vectors each element of $\mathcal{H}_{n,n'}(\mathcal{X})$ can be stored in $O(n' \cdot (\log n + \log |\mathcal{X}|))$ bits, which is shorter than the vector representation when $n' = o(|\mathcal{X}| / \log |\mathcal{X}|)$.

³Note that the order in which bins are released can result in a breach of privacy (e.g. releasing the bins of elements in the dataset before the bins of elements not in the dataset). As a result, our algorithms always sort the released bins according to a predefined ordering based only on \mathcal{X} .

2.2.2 Accuracy

In order to preserve privacy, our algorithms return histograms with noise added to the counts. Therefore, it is crucial to understand their accuracy guarantees. So given a dataset $D \in \mathcal{X}^n$ we compare the **noisy count** $\tilde{c}_x = \mathcal{M}(D)_x$ of $x \in \mathcal{X}$ (the count released by algorithm \mathcal{M}) to its **true count**, $c_x(D)$. We focus on the following two metrics:

Definition 2.5. A histogram algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \mathbb{N}^{\mathcal{X}}$ has (a, β) -**per-query accuracy** if

$$\forall D \in \mathcal{X}^n \quad \forall x \in \mathcal{X} \quad \Pr[|\mathcal{M}(D)_x - c_x(D)| \leq a] \geq 1 - \beta$$

Definition 2.6. A histogram algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \mathbb{N}^{\mathcal{X}}$ has (a, β) -**simultaneous accuracy** if

$$\forall D \in \mathcal{X}^n \quad \Pr[\forall x \in \mathcal{X} \quad |\mathcal{M}(D)_x - c_x(D)| \leq a] \geq 1 - \beta$$

Respectively, these metrics capture the maximum error for any one bin and the maximum error simultaneously over all bins. Even though simultaneous accuracy is commonly used in differential privacy, per-query accuracy has several advantages:

- For histograms, one can achieve a smaller per-query error than is possible for simultaneous error. Indeed, the optimal simultaneous error for $(\epsilon, 0)$ -differentially private histograms is $a = \Theta(\log(|\mathcal{X}|/\beta)/\epsilon)$ whereas the optimal per-query error is $a = \Theta(\log(1/\beta)/\epsilon)$, which is independent of $|\mathcal{X}|$ [HT10, BBKN14].
- Per-query accuracy may be easier to convey to an end user of differential privacy. For example, it is the common interpretation of error bars shown on a graphical depiction of a histogram.

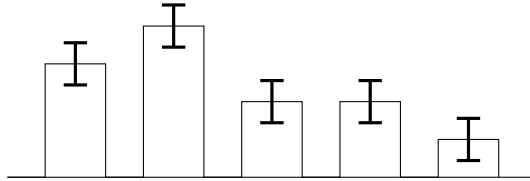


Figure 1: A histogram with error bars

- For many algorithms (such as ours), per-query accuracy is good enough to imply optimal simultaneous accuracy. Indeed, an algorithm with (a, β) -per-query accuracy also achieves $(a, \beta \cdot |\mathcal{X}|)$ -simultaneous accuracy (by a union bound).

However, we may not always be able to achieve as good per-query accuracy as we want. So we will also use the following relaxation which bounds the error only on bins with large enough true count.

Definition 2.7. A histogram algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \mathbb{N}^{\mathcal{X}}$ has (a, β) -**per-query accuracy on counts larger than t** if

$$\forall D \in \mathcal{X}^n \quad \forall x \in \mathcal{X} \text{ s.t. } c_x(D) > t \quad \Pr[|\mathcal{M}(D)_x - c_x(D)| \leq a] \geq 1 - \beta$$

2.3 Probability Terminology

Definition 2.8. Let Z be an integer-valued random variable. The **probability mass function** of Z , denoted f_Z , is the function $f_Z(z) = \Pr[Z = z]$ for all $z \in \mathbb{Z}$. The **cumulative distribution function** of Z , denoted F_Z , is the function $F_Z(z) = \Pr[Z \leq z]$ for all $z \in \mathbb{Z}$. The **support** of Z , denoted $\text{supp}(Z)$, is the set of elements for which $f(z) \neq 0$.

Definition 2.9. Let Y and Z be random variables taking values in discrete range \mathcal{R} . The **statistical distance between Y and Z** (a.k.a. total variation distance) is defined as

$$\begin{aligned} \Delta(Y, Z) &= \max_{A \subseteq \mathcal{R}} |\Pr[Y \in A] - \Pr[Z \in A]| \\ &= \frac{1}{2} \cdot \sum_{a \in \mathcal{R}} |\Pr[Z = a] - \Pr[Y = a]| \end{aligned}$$

Lemma 2.10. Let Y and Z be random variables over discrete range \mathcal{R} . Statistical distance has the following properties:

1. Y and Z are identically distributed, denoted $Y \sim Z$, if and only if $\Delta(Y, Z) = 0$ (equivalently, $F_Y(z) = F_Z(z)$ for all $z \in \mathcal{R}$).
2. Let $T : \mathcal{R} \rightarrow \mathcal{R}'$ be a randomized mapping with \mathcal{R}' discrete. Then

$$\Delta(T(Y), T(Z)) \leq \Delta(Y, Z)$$

3. For $i \in \{1, 2\}$, let Y_i and Z_i be random variables over discrete range \mathcal{R}_i . Then

$$\Delta((Y_1, Y_2), (Z_1, Z_2)) \leq \Delta(Y_1, Z_1) + \max_{a \in \mathcal{R}_1} \Delta(Y_2 | \{Y_1 = a\}, Z_2 | \{Z_1 = a\})$$

Definition 2.11. Let Z_1, \dots, Z_ℓ be integer-valued random variables. The i -th order statistic of Z_1, \dots, Z_ℓ denoted $Z_{(i)}$ is the i -th smallest value among Z_1, \dots, Z_ℓ .

2.3.1 Sampling

Because we are interested in the computational efficiency of our algorithms we need to consider the efficiency of sampling from various distributions.

A standard method for sampling a random variable is via **inverse transform sampling**. Let $\text{Unif}(A)$ denote the uniform distribution over the set A .

Lemma 2.12. Let $U \sim \text{Unif}((0, 1])$. Then for any integer-valued random variable Z we have $F_Z^{-1}(U) \sim Z$ where $F_Z^{-1}(u)$ is defined as $\min\{z \in \text{supp}(Z) : F_Z(z) \geq u\}$.

If Z , the random variable we wish to sample, has finite support we can compute the inverse cumulative distribution by performing binary search on $\text{supp}(Z)$ to find the minimum. This method removes the need to compute the inverse function of the cumulative distribution function. If in addition, the cumulative distribution function of Z can be represented by rational numbers, then we only need to sample from a discrete distribution instead of $(0, 1]$.

Lemma 2.13. Let Z be an integer-valued random variable with finite support and has all probabilities of its cumulative distribution function expressible as rational numbers with denominator d . Then $F_Z^{-1}(U) \sim Z$ where $U \sim (1/d) \cdot \text{Unif}([d]_+)$ and $F_Z^{-1}(u)$ is defined as $\min\{z \in \text{supp}(Z) : F_Z(z) \geq u\}$.

2.4 Model of Computation

We analyze the running time of our algorithms with respect to the *w-bit word RAM model* taking w logarithmic in our input length, namely $w = O(\log n + \log \log |\mathcal{X}|)$. In this model, memory accesses and basic operations (arithmetic, comparisons and logical) on w -bit words are constant time. In addition, we assume the data universe is indexed so we can view $\mathcal{X} = [m]_+$ for some $m \in \mathbb{N}$. Some parameters to our algorithms are rational and we represent rationals by pairs of integers. Some of our algorithms will use numbers that span many words. For ease of notation, we will assume multiplication of two x -bit number is $\tilde{O}(x)$ ([vzGG13] Theorem 8.24).

Our algorithms require randomness so we assume that they have access to an oracle that when given a number $d \in \mathbb{N}_+$ returns a uniformly random integer between 1 and d inclusive.

Finally, for representing histograms as partial vectors, we will assume internally to the algorithms that they are stored as red-black trees. This will allow us to insert and search for elements in $O(\log n \cdot \log |\mathcal{X}|)$ time ([CLRS09] Chapter 13). When releasing a partial vector we don't return the tree itself (which may violate privacy), but instead return a list of bins using an in-order traversal of tree.

3 A General Framework for Implementing Differential Privacy

In this section, we outline a basic framework for implementing a pure differentially private algorithm \mathcal{M} on a finite computer with only a small loss in privacy and possibly a small loss in accuracy. It can be broken down into the following steps:

1. Start by discretizing the input and output of \mathcal{M} so that they can only take on a finite number of values (e.g. rounding a real-valued number to the nearest integer in some finite set). Depending on how utility is measured, the loss in accuracy by discretizing may be acceptable.
2. Then find an algorithm \mathcal{M}' that runs on a finite computer and approximates the output distribution of the discretized version of \mathcal{M} to within “small” statistical distance. Notice that \mathcal{M}' is only guaranteed to satisfy approximate differential privacy and may not satisfy pure differential privacy. (This step may require a non-trivial amount of work. For one example, see Theorem 6.8.)
3. Finally, provided that the statistical distance of the previous step is small enough, by mixing \mathcal{M}' with uniformly random output (from the discretized and finite output space), the resulting algorithm satisfies pure differential privacy.

We will use this framework several times in designing our algorithms. Here we start by formalizing Step 3. That is, for algorithms whose output distribution is close in statistical distance to that of a pure differentially private algorithm, we construct an algorithm satisfying pure differential privacy by mixing it with random output inspired by similar techniques in [KLN⁺11, CDK17].

Algorithm 3.1. $\mathcal{M}_{\mathcal{M}', \mathcal{D}, \gamma}^*(D)$ for $D \in \mathcal{X}^n$ where \mathcal{R} is discrete and finite, an algorithm $\mathcal{M}' : \mathcal{X}^n \rightarrow \mathcal{R}$, a distribution \mathcal{D} over \mathcal{R} and $\gamma \in \mathbb{N}^{-1}$

1. With probability $1 - \gamma$ release $\mathcal{M}'(D)$.
 2. Otherwise release an element sampled from the distribution \mathcal{D} .
-

Lemma 3.2. *Suppose that there is an $(\epsilon, 0)$ -differentially private algorithm $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{R}$ such that $\Delta(\mathcal{M}(D), \mathcal{M}'(D)) \leq \delta$ for all input datasets $D \in \mathcal{X}^n$ with parameter $\delta \in [0, 1)$. Then the algorithm $\mathcal{M}_{\mathcal{M}', \mathcal{D}, \gamma}^* : \mathcal{X}^n \rightarrow \mathcal{R}$ has the following properties:*

i. $(\epsilon, 0)$ -differential privacy whenever

$$\delta \leq \frac{e^\epsilon - 1}{e^\epsilon + 1} \cdot \frac{\gamma}{1 - \gamma} \cdot \min_{r \in \mathcal{R}} \left\{ \Pr_{Z \sim \mathcal{D}}[Z = r] \right\} \quad (1)$$

ii. Running time $O(\log(1/\gamma)) + \text{Time}(\mathcal{M}') + \text{Time}(\mathcal{D})$ where $\text{Time}(\mathcal{D})$ is the time to sample from the distribution \mathcal{D} .

By taking γ and δ small enough and satisfying (1), the algorithm $\mathcal{M}_{\mathcal{M}', \mathcal{D}, \gamma}^*$ satisfies pure differential privacy and has nearly the same utility as \mathcal{M} (due to having a statistical distance at most $\gamma + \delta$ from \mathcal{M}) while allowing for a possibly more efficient implementation since we only need to approximately sample from the output distribution of \mathcal{M} .

To maximize the minimum in (1), one can take $\mathcal{D} \sim \text{Unif}(\mathcal{R})$. However, it may be the case that sampling this distribution exactly is inefficient and we are willing to trade needing a smaller δ to maintain pure differential privacy for a faster sampling algorithm.

Proof of i. Let $r \in \mathcal{R}$ and $p = \Pr_{Z \sim \mathcal{D}}[Z = r]$. Then for neighboring datasets $D, D' \in \mathcal{X}^n$

$$\begin{aligned} \Pr[\mathcal{M}_{\mathcal{M}', \mathcal{D}, \gamma}^*(D) = r] &= \gamma \cdot p + (1 - \gamma) \cdot \Pr[\mathcal{M}'(D) = r] \\ &\leq \gamma \cdot p + (1 - \gamma) \cdot (\Pr[\mathcal{M}(D) = r] + \delta) \\ &\leq \gamma \cdot p + (1 - \gamma) (e^\epsilon \cdot \Pr[\mathcal{M}(D') = r] + \delta) \\ &\leq \gamma \cdot p + (1 - \gamma) (e^\epsilon \cdot (\Pr[\mathcal{M}'(D') = r] + \delta) + \delta) \end{aligned}$$

Rearranging terms and using the upper bound on δ yields

$$\begin{aligned} \Pr[\mathcal{M}_{\mathcal{M}', \mathcal{D}, \gamma}^*(D) = r] &\leq e^\epsilon (1 - \gamma) \cdot \Pr[\mathcal{M}'(D') = r] + \gamma \cdot p + (e^\epsilon + 1) (1 - \gamma) \cdot \delta \\ &\leq e^\epsilon (1 - \gamma) \cdot \Pr[\mathcal{M}'(D') = r] + \gamma \cdot p + (e^\epsilon - 1) \cdot \gamma \cdot \min_{r' \in \mathcal{R}} \left\{ \Pr_{Z \sim \mathcal{D}}[Z = r'] \right\} \\ &\leq e^\epsilon ((1 - \gamma) \cdot \Pr[\mathcal{M}'(D') = r] + \gamma \cdot p) \\ &= e^\epsilon \cdot \Pr[\mathcal{M}_{\mathcal{M}', \mathcal{D}, \gamma}^*(D') = r] \quad \square \end{aligned}$$

Proof of ii. This follows directly from the construction of \mathcal{M}^* . □

4 Counting Queries

Before discussing algorithms for privately releasing histograms, we show how to privately answer a single counting query using only integers of bounded length. While there exist known algorithms for this problem [DKM⁺06, Mir12], our algorithms have additional properties that will be used to construct histogram algorithms in later sections. In general, counting queries have as input the dataset $D \in \mathcal{X}^n$ and the bin x to query. However, we will take the true count, $c_x(D)$, as the input to our counting query algorithms. When constructing histogram algorithms in later sections, this will allow us to improve the running time as we will only need to iterate through the dataset once to determine all true counts prior to answering any counting query. In addition, we would like to keep track of the randomness used by most of our algorithms so we write that as an explicit second input. As a result, we have the following definitions:

Definition 4.1. Let $n, d \in \mathbb{N}_+$. We say a (deterministic or randomized) algorithm $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ is (ε, δ) -**differentially private for counting queries** if the randomized algorithm $\mathcal{M} : \{0, 1\}^n \rightarrow [n]$ defined as $\mathcal{M}(D) = \mathcal{M}(c, U)$ where $c = \sum_{i=1}^n D_i$ and $U \sim \text{Unif}([d]_+)$ is (ε, δ) -differentially private.

Definition 4.2. Let $n, d \in \mathbb{N}_+$. We say $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ has (a, β) -**accuracy** if for all $c \in [n]$

$$\Pr[|\mathcal{M}(c, U) - c| \leq a] \geq 1 - \beta$$

where $U \sim \text{Unif}([d]_+)$.

Definition 4.3. Let $n, d \in \mathbb{N}_+$ and $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be deterministic. Let the **scaled cumulative distribution function of \mathcal{M} at 0** denoted $F_{\mathcal{M}}$ be the function $F_{\mathcal{M}} : [n] \rightarrow [d]_+$ defined as $F_{\mathcal{M}}(z) = d \cdot F_{\mathcal{M}(0, U)}(z)$ where $U \sim \text{Unif}([d]_+)$ for all $z \in [n]$.

Being able to efficiently compute the scaled cumulative distribution function of \mathcal{M} at 0 will be a necessary property for constructing efficient histogram algorithms later (see Proposition 6.19 and Lemma 8.4). Definition 4.1 can easily be extended to handle point queries over data universes with more than two elements.

Lemma 4.4. Let $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be (ε, δ) -differentially private for counting queries. Let $D, D' \in \mathcal{X}^n$ be neighboring datasets. Then for all $x \in \mathcal{X}$ and $c \in [n]$

$$\Pr[\mathcal{M}(c_x(D), U) = c] \leq e^\varepsilon \cdot \Pr[\mathcal{M}(c_x(D'), U) = c] + \delta$$

where $U \sim \text{Unif}([d]_+)$.

Proof. Define the dataset $D^{(x)} \in \{0, 1\}^n$ as

$$D_i^{(x)} = \begin{cases} 1 & \text{if } D_i = x \\ 0 & \text{otherwise} \end{cases}$$

Notice that $c_x(D) = \sum_{i=1}^n D_i^{(x)}$. Similarly, we define the dataset $D'^{(x)}$ for D' . Now, $D^{(x)}$ and $D'^{(x)}$ are neighboring datasets. The lemma follows by (ε, δ) -differential privacy for counting queries. \square

4.1 The Geometric Mechanism

As shown by Dwork, McSherry, Nissim and Smith [DMNS06], we can privately release a counting query by adding appropriately scaled Laplace noise to the count. Because our algorithm's outputs are counts, we do not need to use continuous noise and instead use a discrete analogue, as in [DKM⁺06, GRS12].

We say an integer-valued random variable Z follows a **two-sided geometric distribution with scale parameter s centered at $c \in \mathbb{Z}$** (denoted $Z \sim c + \text{Geo}(s)$) if its probability mass function $f_Z(z)$ is proportional to $e^{-|z-c|/s}$. It can be verified that f_Z and its cumulative distribution function F_Z are

$$f_Z(z) = \left(\frac{e^{1/s} - 1}{e^{1/s} + 1} \right) \cdot e^{-|z-c|/s} \quad F_Z(z) = \begin{cases} \frac{e^{1/s}}{e^{1/s}+1} \cdot e^{-(c-z)/s} & \text{if } z \leq c \\ 1 - \frac{1}{e^{1/s}+1} \cdot e^{-(z-c)/s} & \text{otherwise} \end{cases}$$

for all $z \in \mathbb{Z}$. When c is not specified, it is assumed to be 0. The inverse cumulative distribution of Z is

$$F_Z^{-1}(u) = c + \begin{cases} \lceil s \ln(u) + s \ln(e^{1/s} + 1) \rceil - 1 & \text{if } u \leq 1/2 \\ \lceil -s \ln(1-u) - s \ln(e^{1/s} + 1) \rceil & \text{otherwise} \end{cases}$$

or, equivalently,

$$F_Z^{-1}(u) = c + \left\lceil s \cdot \text{sign}(1/2 - u) \left(\ln(1 - |2u - 1|) + \ln(e^{1/s} + 1) - \ln 2 \right) \right\rceil + \lfloor 2u \rfloor - 1$$

Now, we state the counting query algorithm using discrete noise formally studied in [GRS12]. We will not keep track of the randomness used by this algorithm, but to match our syntax for counting query algorithms we use the dummy parameter 1 as the second argument.

Algorithm 4.5. $\text{GeometricMechanism}_{n,\varepsilon}(c, 1)$ for $c \in [n]$ where $n \in \mathbb{N}_+$ and $\varepsilon > 0$

1. Return \tilde{c} set to $c + \text{Geo}(2/\varepsilon)$ clamped to the interval $[0, n]$. i.e.

$$\tilde{c} = \begin{cases} 0 & \text{if } Z \leq 0 \\ n & \text{if } Z \geq n \\ Z & \text{otherwise} \end{cases} \quad \text{where } Z = c + \text{Geo}(2/\varepsilon).$$

Theorem 4.6. *Let $n \in \mathbb{N}_+$ and $\varepsilon > 0$. Then $\text{GeometricMechanism}_{n,\varepsilon} : [n] \times [1]_+ \rightarrow [n]$ has the following properties:*

- i. $\text{GeometricMechanism}_{n,\varepsilon}$ is $(\varepsilon/2, 0)$ -differentially private for counting queries [GRS12].
- ii. $\text{GeometricMechanism}_{n,\varepsilon}$ has (a, β) -accuracy for $\beta \in (0, 1]$ and

$$a = \left\lceil \frac{2}{\varepsilon} \cdot \ln \frac{1}{\beta} \right\rceil$$

Proof of ii. Let $Z \sim \text{Geo}(2/\varepsilon)$. Then for $c \in [n]$,

$$\begin{aligned}
\Pr[|\text{GeometricMechanism}_{n,\varepsilon}(c) - c| \leq a] &\geq \Pr[|Z| \leq \lfloor a \rfloor] \\
&= 1 - 2 \cdot \Pr[Z \leq -\lfloor a \rfloor - 1] \\
&= 1 - 2 \cdot \frac{e^{-\lfloor a \rfloor \cdot \varepsilon/2}}{e^{\varepsilon/2} + 1} \\
&\geq 1 - \frac{2 \cdot \beta}{e^{\varepsilon/2} + 1} \\
&\geq 1 - \beta \quad \square
\end{aligned}$$

As presented above, this algorithm needs to store integers of unbounded size since $\text{Geo}(2/\varepsilon)$ is unbounded in magnitude. As noted in [GRS12], by restricting the generated noise to a fixed range we can avoid this problem. However, even when the generated noise is restricted to a fixed range, generating this noise via inverse transform sampling may require infinite precision. By appropriately choosing ε , the probabilities of this noise's cumulative distribution function can be represented with finite precision, and therefore generating this noise via inverse transform sampling only requires finite precision.

Theorem 4.7. *Let $n \in \mathbb{N}_+$, $\varepsilon \in \mathbb{N}^{-1}$ and $\tilde{\varepsilon} = 2 \cdot \ln(1 + 2^{-\lceil \log(2/\varepsilon) \rceil}) \in (4/9 \cdot \varepsilon, \varepsilon]$. Then there is a deterministic algorithm $\text{GeoSample}_{n,\varepsilon} : [n] \times [d]_+ \rightarrow [n]$ where $\log d = O(n \cdot \log(1/\varepsilon))$ with the following properties:*

- i. $\text{GeoSample}_{n,\varepsilon}(c, U) \sim \text{GeometricMechanism}_{n,\tilde{\varepsilon}}(c, 1)$ for all $c \in [n]$ where $U \sim \text{Unif}([d]_+)$. Thus, $\text{GeoSample}_{n,\varepsilon}$ is $(\tilde{\varepsilon}/2, 0)$ -differentially private for counting queries and has (a, β) -accuracy for $\beta \in (0, 1]$ and*

$$a = \left\lceil \frac{2}{\tilde{\varepsilon}} \cdot \ln \frac{1}{\beta} \right\rceil$$

- ii. $\text{GeoSample}_{n,\varepsilon}$ has running time $\tilde{O}(n \cdot \log(1/\varepsilon))$.*
- iii. For all $z \in [n]$, $F_{\text{GeoSample}_{n,\varepsilon}}(z)$ can be computed in time $\tilde{O}(n \cdot \log(1/\varepsilon))$.*
- iv. $\text{GeoSample}_{n,\varepsilon}(c, u)$ is a non-decreasing function in u .*

We have chosen $\tilde{\varepsilon}$ so that the cumulative distribution function of a two-sided geometric random variable with scale parameter $2/\tilde{\varepsilon}$ clamped to $[0, n]$ takes on only rational values with a common denominator d . Therefore, to implement inverse transform sampling on this distribution we only need to choose a uniformly random integer from $[d]_+$ rather than a uniformly random variable over $(0, 1]$ which allows us to provide a strict bound on the running time.

Algorithm 4.8. $\text{GeoSample}_{n,\varepsilon}(c, u)$ for $c \in [n]$ and $u \in [d]_+$ where $n \in \mathbb{N}_+$ and $\varepsilon \in \mathbb{N}^{-1}$

1. Let $k = \lceil \log(2/\varepsilon) \rceil$ and $d = (2^{k+1} + 1)(2^k + 1)^{n-1}$.
2. For $z \in \mathbb{Z}$, define the function

$$F(z) = \begin{cases} 0 & \text{if } z < 0 \\ 2^{k(c-z)} (2^k + 1)^{n-(c-z)} & \text{if } z \in [0, c) \\ d - 2^{k(z-c+1)} (2^k + 1)^{n-1-(z-c)} & \text{if } z \in [c, n) \\ d & \text{if } z \geq n \end{cases}$$

3. Using binary search find the smallest $z \in [n]$ such that $F(z) \geq u$.
 4. Return z .
-

The function F is obtained by clearing denominators in the cumulative distribution function of $c + \text{Geo}(2/\tilde{\varepsilon})$ clamped to $[0, n]$.

Lemma 4.9. *Let $\tilde{\varepsilon}$, c , k , d and F be defined as in Theorem 4.7 and Algorithm 4.8. Then $F(z) \in [d]$ and $F(z)/d$ equals the cumulative distribution function of $c + \text{Geo}(2/\tilde{\varepsilon})$ clamped to $[0, n]$.*

We prove this lemma after seeing how it implies Theorem 4.7.

Proof of Theorem 4.7 Part i. By construction, for all $z \in [n]$

$$\Pr[\text{GeoSample}_{n,\varepsilon}(c, U) \leq z] = \Pr[U \leq F(z)] = F(z)/d$$

implying $\text{GeoSample}_{n,\varepsilon}(c, U) \sim \text{GeometricMechanism}_{n,\tilde{\varepsilon}}(c, 1)$ by Lemma 4.9. \square

Proof of Theorem 4.7 Part ii-iv. Notice that integers used do not exceed d whose bit length is $O(n \cdot \log(1/\varepsilon))$. Thus, $F(z)$ can be computed in $\tilde{O}(n \cdot \log(1/\varepsilon))$ time using exponentiation by repeated squaring. By construction, $F(z) = F_{\text{GeoSample}_{n,\varepsilon}}(z)$, when $c = 0$ implying Part *iii*.

Notice that the binary search of $\text{GeoSample}_{n,\varepsilon}$ has at most $O(\log n)$ rounds each with an evaluation of F . Thus, $\text{GeoSample}_{n,\varepsilon}$ has the desired running time and by construction $\text{GeoSample}_{n,\varepsilon}(c, u)$ is a non-decreasing function in u . \square

Proof of Lemma 4.9. The cumulative distribution function of $Z \sim c + \text{Geo}(2/\tilde{\varepsilon})$ is

$$F_Z(z) = \begin{cases} 0 & \text{if } z < 0 \\ \frac{e^{\tilde{\varepsilon}/2}}{e^{\tilde{\varepsilon}/2} + 1} \cdot e^{-(c-z) \cdot \tilde{\varepsilon}/2} & \text{if } z \in [0, c) \\ 1 - \frac{1}{e^{\tilde{\varepsilon}/2} + 1} \cdot e^{-(z-c) \cdot \tilde{\varepsilon}/2} & \text{if } z \in [c, n) \\ 1 & \text{if } z \geq n \end{cases}$$

Consider the case when $z \in [0, c)$.

$$\begin{aligned}
F_Z(z) &= \frac{e^{\tilde{\varepsilon}/2}}{e^{\tilde{\varepsilon}/2} + 1} \cdot e^{-(c-z)\tilde{\varepsilon}/2} = \frac{1 + 2^{-k}}{2 + 2^{-k}} \cdot (1 + 2^{-k})^{-(c-z)} \\
&= \frac{2^k + 1}{2^{k+1} + 1} \cdot \left(\frac{2^k}{2^k + 1} \right)^{c-z} \\
&= \frac{2^{k(c-z)}}{(2^{k+1} + 1)(2^k + 1)^{c-z-1}} \cdot \left(\frac{2^k + 1}{2^k + 1} \right)^{n-(c-z)} \\
&= \frac{2^{k(c-z)} (2^k + 1)^{n+z-c}}{d} \\
&= \frac{F(z)}{d}
\end{aligned}$$

A similar argument holds for $z \in [c, n)$. The remaining cases are trivial. So $F_Z(z) = F(z)/d$ for all $z \in \mathbb{Z}$. \square

4.2 Approximating Geometric Noise to Release Counting Queries Faster

Notice that $\text{GeoSample}_{n,\varepsilon}$ has running time at least linear in n . This is due to evaluating a (scaled) cumulative distribution function operating on integers with bit length $\Omega(n)$. We can improve the running time by approximately sampling from a two-sided geometric distribution. Small tail probabilities are dropped to reduce the number of required bits to represent probabilities to logarithmic in n . And then to recover pure differential privacy, following Lemma 3.2, we mix with uniformly random output.

Theorem 4.10. *Let $n \in \mathbb{N}_+$, $\varepsilon, \gamma \in \mathbb{N}^{-1}$ and $\tilde{\varepsilon} = 2 \cdot \ln(1 + 2^{-\lceil \log(2/\varepsilon) \rceil}) \in (4/9 \cdot \varepsilon, \varepsilon]$. Then there is a deterministic algorithm $\text{FastSample}_{n,\varepsilon,\gamma} : [n] \times [d]_+ \rightarrow [n]$ where $\log d = \tilde{O}(1/\varepsilon) \cdot \log(n/\gamma)$ with the following properties:*

- i. $\text{FastSample}_{n,\varepsilon,\gamma}$ is $(\varepsilon/2, 0)$ -differentially private for counting queries.
- ii. For every $\beta > \gamma$, $\text{FastSample}_{n,\varepsilon,\gamma}$ has (a, β) -accuracy for

$$a = \left\lceil \frac{2}{\tilde{\varepsilon}} \ln \frac{1}{\beta - \gamma} \right\rceil$$

- iii. $\text{FastSample}_{n,\varepsilon,\gamma}$ has running time

$$\tilde{O}\left(\frac{1}{\varepsilon} \cdot \log^2 n + \frac{1}{\varepsilon} \cdot \log n \cdot \log \frac{1}{\gamma}\right)$$

- iv. For all $z \in [n]$, $F_{\text{FastSample}_{n,\varepsilon,\gamma}}(z)$ can be computed in time

$$\tilde{O}\left(\frac{1}{\varepsilon} \cdot \log \frac{n}{\gamma}\right)$$

- v. $\text{FastSample}_{n,\varepsilon,\gamma}(c, u)$ is a non-decreasing function in u .

Algorithm 4.11. $\text{FastSample}_{n,\varepsilon,\gamma}(c, u)$ for $c \in [n]$ and $u \in [d]_+$ where $n \in \mathbb{N}_+$ and $\varepsilon, \gamma \in \mathbb{N}^{-1}$

1. Let $k = \lceil \log(2/\varepsilon) \rceil$ and $t = \left\lceil \frac{9}{2\varepsilon} \cdot \left\lceil \log \left(\frac{8(n+1)(1-\gamma)}{\varepsilon\gamma} \right) \right\rceil \right\rceil - 1$.
2. Let $d' = (2^{k+1} + 1)(2^k + 1)^t$ and $d = (n + 1) \cdot d'/\gamma$.
3. For $z \in [n]$, define the functions

$$F'(z) = \begin{cases} 0 & \text{if } z < \max\{0, c - t\} \\ 2^{k(c-z)}(2^k + 1)^{t+1-(c-z)} - 2^{k(t+1)} & \text{if } z \in [\max\{0, c - t\}, c) \\ d' - 2^{k(z-c+1)}(2^k + 1)^{t-(z-c)} + 2^{k(t+1)} & \text{if } z \in [c, \min\{c + t, n\}) \\ d' & \text{if } z \geq \min\{c + t, n\} \end{cases}$$

and

$$F(z) = (z + 1) \cdot d' + (1/\gamma - 1) \cdot (n + 1) \cdot F'(z)$$

4. Using binary search find the smallest $z \in [n]$ such that $F(z) \geq u$.
 5. Release z .
-

As presented, it is clear how to compute the cumulative distribution function of this algorithm's output distribution, a necessary property for Section 6. And the algorithm, as a function of u , is non-decreasing, a property that will be used in Section 8. However, as stated, the interpretation of F' and F may not be clear. This is clarified by the following lemma and Figure 2.

Lemma 4.12. *Let $k, t, d, d', F'(z), F(z)$ be defined as in Algorithm 4.11. Then*

- i. $F'(z) \in [d']$ and $F'(z)/d'$ is the cumulative distribution function of a random variable Z' clamped to $[0, n]$ where Z' has probability mass function*

$$f_{Z'}(z) = \begin{cases} \Pr[Z = z] & \text{if } z \in [c - t, c + t] \text{ and } z \neq c \\ \Pr[Z = c] + \Pr[|Z - c| > t] & \text{if } z = c \\ 0 & \text{if } z \notin [c - t, c + t] \end{cases}$$

for all $z \in \mathbb{Z}$ where $Z \sim c + \text{Geo}(2/\varepsilon)$.

- ii. $F(z) \in [d]$ and $F(z)/d$ is the cumulative distribution function of the random variable that with probability $1 - \gamma$ is distributed as Z' (defined in Part i) clamped to $[0, n]$ and with probability γ is uniform over $[n]$.*

We prove this lemma after using it to prove Theorem 4.10.

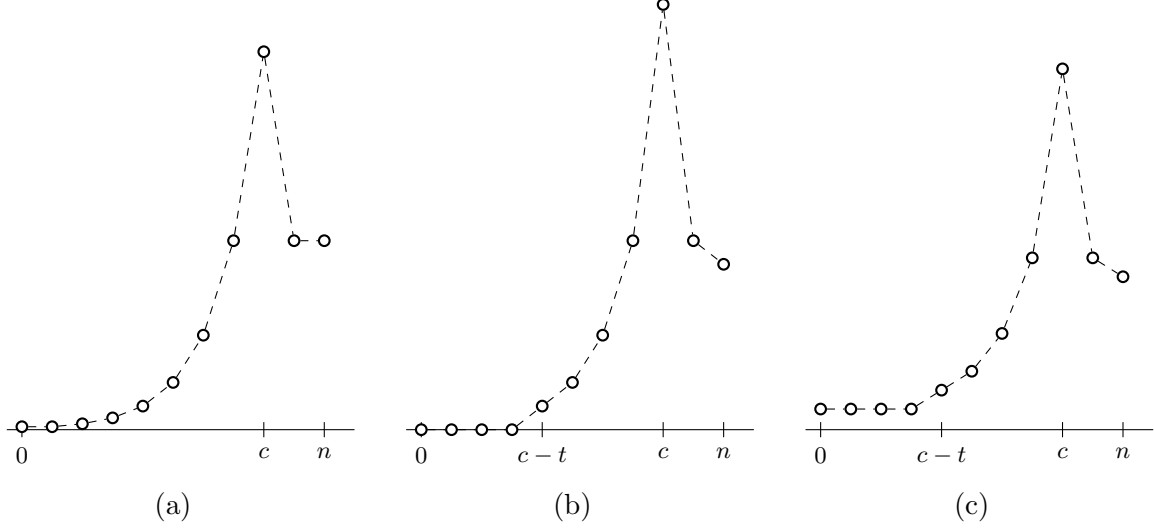


Figure 2: Example probability mass functions of (a) **GeoSample**, (b) Z' clamped to $[0, n]$ as defined in Lemma 4.12 and (c) **FastSample**. Notice that in this example $c + t > n$.

Proof of Theorem 4.10 Part i. By construction, for all $z \in [n]$

$$\Pr[\mathbf{FastSample}_{n,\varepsilon,\gamma}(c,U) \leq z] = \Pr[U \leq F(z)]$$

where $U \sim \text{Unif}([d]_+)$. So $\mathbf{FastSample}_{n,\varepsilon,\gamma}(c,U)$ has cumulative distribution function $F(z)/d$. Then by Lemma 4.12, $\mathbf{FastSample}_{n,\varepsilon,\gamma}(c,U)$ is a mixture that with probability $1 - \gamma$ is distributed as Z' clamped to $[0, n]$ as defined in Lemma 4.12 and otherwise is distributed as $\text{Unif}([n])$.

Since $\mathbf{GeoSample}_{n,\varepsilon}$ is $(\varepsilon/2, 0)$ -differentially private for counting queries (Theorem 4.7), we can use Lemma 3.2 to prove that $\mathbf{FastSample}_{n,\varepsilon,\gamma}$ is also $(\varepsilon/2, 0)$ -differentially private for counting queries provided we can show the statistical distance between Z' clamped to $[0, n]$ and $\mathbf{GeoSample}_{n,\varepsilon}$ is small enough.

$$\begin{aligned} \Delta(Z' \text{ clamped to } [0, n], \mathbf{GeoSample}_{n,\varepsilon}(c,U)) &\leq \Delta(Z', Z) \\ &= \Pr[|Z - c| > t] \\ &= 2 \cdot F_Z(c - t - 1) \\ &= 2 \cdot \frac{e^{\tilde{\varepsilon}/2}}{e^{\tilde{\varepsilon}/2} + 1} \cdot e^{-\tilde{\varepsilon} \cdot (t+1)/2} \\ &\leq \frac{4}{3} \cdot e^{-\frac{\tilde{\varepsilon}}{2} \cdot \frac{9}{2\varepsilon} \cdot \ln\left(\frac{8(n+1)(1-\gamma)}{\varepsilon\gamma}\right)} \\ &\leq \frac{\varepsilon}{6} \cdot \frac{\gamma}{1-\gamma} \cdot \frac{1}{n+1} \\ &\leq \frac{e^{\varepsilon/2} - 1}{e^{\varepsilon/2} + 1} \cdot \frac{\gamma}{1-\gamma} \cdot \frac{1}{n+1} \end{aligned}$$

where $Z \sim \text{Geo}(2/\tilde{\varepsilon})$. □

Proof of Theorem 4.10 Part ii. Let $U \sim \text{Unif}([d]_+)$. Let Z' be defined as in Lemma 4.12 and

$Z \sim c + \text{Geo}(2/\tilde{\varepsilon})$. Then

$$\begin{aligned} \Pr[|\text{FastSample}_{n,\varepsilon,\gamma}(c,U) - c| \leq a] &\geq (1 - \gamma) \cdot \Pr[|Z' - c| \leq a] \\ &\geq (1 - \gamma) \cdot \Pr[|Z - c| \leq a] \\ &\geq (1 - \gamma) \cdot (1 - (\beta - \gamma)) \\ &\geq 1 - \beta \end{aligned} \quad \square$$

Proof of Theorem 4.10 Part iii-v. d' can be computed in $\tilde{O}(t \cdot \log(1/\varepsilon))$ time using exponentiation by repeated squaring. Notice that integers used in computing F' and F do not exceed d whose bit length is $O(t \cdot \log(1/\varepsilon) + \log(n/\gamma))$. Thus, $F(z)$ can be computed in time

$$\tilde{O}\left(t \cdot \log \frac{1}{\varepsilon} + \log \frac{n}{\gamma}\right)$$

Part *iv* follows after observing $t = \tilde{O}(1/\varepsilon) \cdot \log(n/\gamma)$.

Now, the binary search of $\text{FastSample}_{n,\varepsilon,\gamma}$ has at most $O(\log n)$ rounds each with an evaluation of $F(z)$ for some $z \in [n]$. Thus, we obtain the desired running time and by construction $\text{FastSample}_{n,\varepsilon,\gamma}(c,u)$ is a non-decreasing function in u . \square

Proof of Lemma 4.12 Part i. Let Z'' be defined as Z' clamped to $[0, n]$. The cumulative distribution function of Z'' is

$$F_{Z''}(z) = \begin{cases} 0 & \text{if } z < \max\{0, c - t\} \\ F_Z(z) - F_Z(c - t - 1) & \text{if } z \in [\max\{0, c - t\}, c) \\ F_Z(z) + (1 - F_Z(c + t)) & \text{if } z \in [c, \min\{c + t, n\}) \\ 1 & \text{if } z \geq \min\{c + t, n\} \end{cases}$$

where $Z \sim c + \text{Geo}(2/\tilde{\varepsilon})$. Consider the case $z \in [\max\{0, c - t\}, c)$.

$$\begin{aligned} F_{Z''}(z) &= F_Z(z) - F_Z(c - t - 1) \\ &= \left(\frac{2^k + 1}{2^{k+1} + 1}\right) \left(\frac{2^k}{2^k + 1}\right)^{c-z} - \left(\frac{2^k + 1}{2^{k+1} + 1}\right) \left(\frac{2^k}{2^k + 1}\right)^{t+1} \\ &= \frac{1}{d'} \cdot \left(2^{k(c-z)}(2^k + 1)^{t+1-(c-z)} - 2^{k(t+1)}\right) \\ &= \frac{F'(z)}{d'} \end{aligned}$$

and $F'(z) \in [d']$ since $c - z < t + 1$. A similar argument holds for $z \in [c, \min\{c + t, n\})$. The remaining cases are trivial. So $F_{Z''}(z) = F'(z)/d'$ for all $z \in \mathbb{Z}$. \square

Proof of Lemma 4.12 Part ii. Following from Part *i*, notice that for $z \in [n]$

$$\frac{F(z)}{d} = \gamma \cdot \frac{z + 1}{n + 1} + (1 - \gamma) \cdot \frac{F'(z)}{d'} = \gamma \cdot F_{U_{[n]}}(z) + (1 - \gamma) \cdot \frac{F'(z)}{d'}$$

where $U_{[n]} \sim \text{Unif}([n])$ which implies the desired result. \square

5 Generalizations of Known Histogram Algorithms

In this section we show how to construct differentially private histograms within our finite model of computation given a private algorithm for releasing a single counting query.

5.1 The Laplace Mechanism

As shown by Dwork, McSherry, Nissim and Smith [DMNS06], we can privately release a histogram by adding independent and appropriately scaled Laplace noise to each bin. Below we state a generalization guaranteeing privacy provided the counting query algorithm used is private and the released counts are independent.

Algorithm 5.1. $\text{BasicHistogram}_{\mathcal{M},A}(D)$ for $D \in \mathcal{X}^n$ where $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ and $A \subseteq \mathcal{X}$

1. Compute $c_x(D)$ for all $x \in A$.
 2. For each $x \in A$, do the following:
 - (a) Sample u_x uniformly at random from $[d]_+$.
 - (b) Let $\tilde{c}_x = \mathcal{M}(c_x(D), u_x)$.
 - (c) Release (x, \tilde{c}_x) .
-

Note that the output of this algorithm is a collection of bins (x, \tilde{c}_x) representing a partial vector.

Theorem 5.2. *Let $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be $(\varepsilon/2, 0)$ -differentially private for counting queries and have (a, β) -accuracy. And let $A \subseteq \mathcal{X}$. Then $\text{BasicHistogram}_{\mathcal{M},A} : \mathcal{X}^n \rightarrow \mathbb{N}^A$ has the following properties:*

i. $\text{BasicHistogram}_{\mathcal{M},A}$ is $(\varepsilon, 0)$ -differentially private.

ii. For all $D \in \mathcal{X}^n$, we have

$$\forall x \in A \quad \Pr[|(\text{BasicHistogram}_{\mathcal{M},A}(D))_x - c_x(D)| \leq a] \geq 1 - \beta$$

In particular, $\text{BasicHistogram}_{\mathcal{M},\mathcal{X}}(D)$ has (a, β) -per-query accuracy.

iii. For all $D \in \mathcal{X}^n$, we have

$$\Pr[\forall x \in A \quad |(\text{BasicHistogram}_{\mathcal{M},A}(D))_x - c_x(D)| \leq a] \geq 1 - \beta'$$

where $\beta' = 1 - (1 - \beta)^{|A|} \leq \beta \cdot |A|$. In particular, $\text{BasicHistogram}_{\mathcal{M},\mathcal{X}}$ has (a, β') -simultaneous accuracy where $\beta' = 1 - (1 - \beta)^{|\mathcal{X}|} \leq \beta \cdot |\mathcal{X}|$.

iv. $\text{BasicHistogram}_{\mathcal{M},A}$ has running time

$$O(n \log n \cdot \log |\mathcal{X}|) + |A| \cdot O(\log n \cdot \log |\mathcal{X}| + \log d + \text{Time}(\mathcal{M}))$$

It is important to note that the privacy guarantee only holds when A is fixed and does not depend on the dataset D . The choice of parameterizing by A will be convenient in defining more complex histogram algorithms later.

Proof of i. This proof follows similarly to the proof of privacy for the Laplace Mechanism. Let $D, D' \in \mathcal{X}^n$ be neighboring datasets and let $h \in \mathbb{N}^A$. Then

$$\Pr[\forall x \in A \quad \tilde{c}_x = h_x] = \prod_{x \in A} \Pr[\mathcal{M}(c_x(D), u_x) = h_x] \quad \text{by independence}$$

Because there are at most two $x \in A$ for which $c_x(D) \neq c_x(D')$, by the $(\varepsilon/2, 0)$ -differential privacy for counting queries of \mathcal{M} with Lemma 4.4 and composition (Lemma 2.4),

$$\begin{aligned} \prod_{x \in A} \Pr[\mathcal{M}(c_x(D), u_x) = h_x] &\leq e^\varepsilon \cdot \prod_{x \in A} \Pr[\mathcal{M}(c_x(D'), u_x) = h_x] \\ &\leq e^\varepsilon \cdot \Pr[\forall x \in A \ \tilde{c}_x(D') = h_x] \quad \square \end{aligned}$$

Proof of ii. For all $x \in \mathcal{X}$, \tilde{c}_x is distributed as $\mathcal{M}(c_x(D), u_x)$ where $u_x \sim \text{Unif}([d]_+)$. The result follows from \mathcal{M} having (a, β) -accuracy. \square

Proof of iii. Simultaneous accuracy follows similarly as the bins' counts are independent. So

$$\Pr[\forall x \in A \ |\tilde{c}_x - c_x(D)| \leq a] = \prod_{x \in A} \Pr[|\tilde{c}_x - c_x(D)| \leq a] \geq (1 - \beta)^{|A|} \quad \square$$

Proof of iv. Computing $c_x(D)$ for all $x \in \mathcal{X}$ can be accomplished by iterating through the dataset once and maintaining a partial vector with counts for the observed data elements. This can be done in $O(n \log n \cdot \log |\mathcal{X}|)$ time. Each of the $|A|$ bin releases takes time $O(\log n \cdot \log |\mathcal{X}| + \log d + \text{Time}(\mathcal{M}))$ in order to get the true count, generate the randomness and then compute the noisy count. \square

Now, we can use the counting query algorithms of Section 4 to get an explicit instantiation of this algorithm. If we take $\mathcal{M} = \text{GeometricMechanism}$, then $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ is identically distributed to the Truncated Geometric Mechanism of Ghosh, Roughgarden and Sundararajan [GRS12] which achieves per-query and simultaneous accuracy with error up to constant factors matching known lower bounds for releasing a private histogram [HT10, BBKN14].

Theorem 5.3. *Let $\varepsilon, \beta_0 \in \mathbb{N}^{-1}$ and $\mathcal{M} = \text{FastSample}_{n, \varepsilon, \gamma}$ where $\gamma = \beta_0 / (2|\mathcal{X}|)$. Then*

i. $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ is $(\varepsilon, 0)$ -differentially private.

ii. For every $\beta \geq 2\gamma$, $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ has (a, β) -per-query accuracy for

$$a = \left\lceil \frac{9}{2\varepsilon} \ln \left(\frac{2}{\beta} \right) \right\rceil$$

iii. For every $\beta \geq \beta_0$, $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ has (a, β) -simultaneous accuracy for

$$a = \left\lceil \frac{9}{2\varepsilon} \ln \left(\frac{2 \cdot |\mathcal{X}|}{\beta} \right) \right\rceil$$

iv. $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ has running time

$$\tilde{O} \left(\frac{|\mathcal{X}|}{\varepsilon} \cdot \left(\log^2 n + \log n \cdot \log \frac{1}{\beta_0} \right) \right) + O(n \log n \cdot \log |\mathcal{X}|)$$

\mathcal{M}	Running Time	(a, β) -Per-Query	(a, β) -Simul.
GeometricMechanism	n/a	$\left\lceil \frac{2}{\varepsilon} \ln \frac{1}{\beta} \right\rceil$	$\left\lceil \frac{2}{\varepsilon} \ln \frac{ \mathcal{X} }{\beta} \right\rceil$
GeoSample	$\tilde{O}(\mathcal{X} \cdot n \cdot \log(1/\varepsilon))$	$\left\lceil \frac{9}{2\varepsilon} \ln \frac{1}{\beta} \right\rceil$	$\left\lceil \frac{9}{2\varepsilon} \ln \frac{ \mathcal{X} }{\beta} \right\rceil$
FastSample	$\tilde{O}((\mathcal{X} /\varepsilon) \cdot \log^2 n) + \tilde{O}(n) \cdot \log \mathcal{X} $	$\left\lceil \frac{9}{2\varepsilon} \ln \frac{2}{\beta} \right\rceil$	$\left\lceil \frac{9}{2\varepsilon} \ln \frac{2 \mathcal{X} }{\beta} \right\rceil$

Figure 3: The running time and errors of $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ for the counting query algorithms of Section 4. Values shown are for a $(\varepsilon, 0)$ -differentially private release where $\varepsilon \in \mathbb{N}^{-1}$. For FastSample , we assume $\beta_0 \geq 1/n^{O(1)}$ and γ is specified in Theorem 5.3.

5.2 Stability-Based Histogram

For a large data universe \mathcal{X} , the at least linear in $|\mathcal{X}|$ running time of $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ can be prohibitive. By using approximate differential privacy, we can release counts for a smaller number of bins (at most n) based on stability techniques [KKMN09, BNS16]. We present a generalization of the algorithm from [BNS16].

Algorithm 5.4. $\text{StabilityHistogram}_{\mathcal{M}, b}(D)$ for $D \in \mathcal{X}^n$ where $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ and $b \in [n]$

1. Let $A = \{x \in \mathcal{X} : c_x(D) > 0\}$.
 2. Let $\{(x, \tilde{c}_x)\}_{x \in A} = \text{BasicHistogram}_{\mathcal{M}, A}(D)$.
 3. Release $h = \{(x, \tilde{c}_x) : x \in A \text{ and } \tilde{c}_x > b\} \in \mathcal{H}_{n, n}(\mathcal{X})$.
-

Note that we only release counts for $x \in \mathcal{X}$ whose true count is nonzero, namely elements in the set A . Thus, the output length is $O(n \cdot (\log |\mathcal{X}| + \log n))$. However, releasing the set A does not satisfy pure differential privacy because this would distinguish between neighboring datasets: one with a count of 0 and the other with a count of 1 for some element $x \in \mathcal{X}$. Thus, we only release noisy counts \tilde{c}_x that exceed a threshold b . If b is large enough, then a count of 1 will only be kept with small probability, yielding approximate differential privacy.

Theorem 5.5. *Let $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be $(\varepsilon/2, 0)$ -differentially private for counting queries and have (a, β) -accuracy. And let $b \in [n]$. Then $\text{StabilityHistogram}_{\mathcal{M}, b} : \mathcal{X}^n \rightarrow \mathcal{H}_{n, n}(\mathcal{X})$ has the following properties:*

- i. $\text{StabilityHistogram}_{\mathcal{M}, b}$ is (ε, δ) -differentially private whenever*

$$\delta \geq 2 \cdot \Pr[\mathcal{M}(1, U) > b] \text{ for } U \sim \text{Unif}([d]_+)$$

- ii. $\text{StabilityHistogram}_{\mathcal{M}, b}$ has (a, β) -per-query accuracy on counts larger than $a + b$.*

- iii. $\text{StabilityHistogram}_{\mathcal{M}, b}$ has $(a + b, \beta')$ -simultaneous accuracy where*

$$\beta' = 1 - (1 - \beta)^n \leq n \cdot \beta$$

iv. $\text{StabilityHistogram}_{\mathcal{M},b}$ has running time

$$O(n \log n \cdot \log |\mathcal{X}| + n \cdot \log d + n \cdot \text{Time}(\mathcal{M}))$$

Proof of i. Let $D, D' \in \mathcal{X}^n$ be neighboring datasets, $h \sim \text{StabilityHistogram}_{\mathcal{M},b}(D)$ and $h' \sim \text{StabilityHistogram}_{\mathcal{M},b}(D')$. Let $U \sim \text{Unif}([d]_+)$. Let $x \in \mathcal{X}$ such that $c_x(D) \neq c_x(D')$ and let $S \subseteq [n]$. There are 3 cases to consider:

- $c_x(D) \geq 1$ and $c_x(D') \geq 1$. By the privacy guarantee on \mathcal{M} , we have $\Pr[\mathcal{M}(c_x(D), U) \in S] \leq e^{\varepsilon/2} \cdot \Pr[\mathcal{M}(c_x(D'), U) \in S]$. Thus, by differential privacy's closure under post-processing (Lemma 2.2),

$$\Pr[h_x \in S] \leq e^{\varepsilon/2} \cdot \Pr[h'_x \in S]$$

- $c_x(D) = 1$ and $c_x(D') = 0$. Notice that $\Pr[h'_x = 0] = 1$. So if $0 \in S$, then $\Pr[h_x \in S] \leq \Pr[h'_x \in S]$. If $0 \notin S$, then

$$\begin{aligned} \Pr[h_x \in S] &\leq \Pr[h_x > 0] = \Pr[h_x > b] \\ &= \Pr[\mathcal{M}(1, U) > b] \\ &\leq \delta/2 = \Pr[h'_x \in S] + \delta/2 \end{aligned}$$

- $c_x(D) = 0$ and $c_x(D') = 1$. This case follows similarly to the previous one.

Then overall

$$\Pr[h_x \in S] \leq e^{\varepsilon/2} \cdot \Pr[h'_x \in S] + \delta/2$$

Because there are at most two bins on which D and D' have differing counts and each count \tilde{c}_x is computed independently, by Lemma 2.4, this algorithm is (ε, δ) -differentially private. \square

Proof of ii. Let $U \sim \text{Unif}([d]_+)$ and $x \in \mathcal{X}$ such that $c_x(D) > a + b$. Notice that $|\tilde{c}_x - c_x(D)| \leq a$ implies $\tilde{c}_x \geq c_x(D) - a > b$. Thus,

$$\begin{aligned} \Pr[|h_x - c_x(D)| \leq a] &= \Pr[|\tilde{c}_x - c_x(D)| \leq a] \\ &= \Pr[|\mathcal{M}(c_x(D), U) - c_x(D)| \leq a] \\ &\geq 1 - \beta \end{aligned} \quad \square$$

Proof of iii. Notice that the counts of elements not in A are trivially accurate. Therefore, we only need to consider the counts of elements in A . By Theorem 5.2 Part *iii*, $\text{BasicHistogram}_{\mathcal{M},A}(D)$ has (a, β') -simultaneous accuracy where $\beta' = 1 - (1 - \beta)^n$ as $|A| \leq n$.

The final step of $\text{StabilityHistogram}_{\mathcal{M},b}$ can increase the error on any count additively by at most b . Therefore, $\text{StabilityHistogram}_{\mathcal{M},b}$ has $(a + b, \beta')$ -simultaneous accuracy. \square

Proof of iv. The running time follows from Theorem 5.2 Part *iv* where $|A| \leq n$ after noting the final step has running time $O(n \log n \cdot \log |\mathcal{X}|)$. \square

By taking $\mathcal{M} = \text{GeoSample}$, the simultaneous accuracy of $\text{StabilityHistogram}_{\mathcal{M},b}$ matches known lower bounds. However, we only achieve optimal per-query accuracy on sufficiently large counts. This constraint is necessary for any algorithm outputting sparse histograms as we will show in Theorem 7.2. The following theorem shows the accuracies achieved by taking $\mathcal{M} = \text{FastSample}$.

Theorem 5.6. Let $\varepsilon, \delta, \beta_0 \in \mathbb{N}^{-1}$, $\mathcal{M} = \text{FastSample}_{n, \varepsilon, \gamma}$ where $\gamma = \min\{\beta_0/(2n), \delta/4\}$ and $b = 1 + \lceil 9/(2\varepsilon) \cdot \ln(4/\delta) \rceil$. Then

i. $\text{StabilityHistogram}_{\mathcal{M}, b}$ is (ε, δ) -differentially private.

ii. For every $\beta \geq 2\gamma$, $\text{StabilityHistogram}_{\mathcal{M}, b}$ has (a, β) -per-query accuracy on counts larger than t for

$$a = \left\lceil \frac{9}{2\varepsilon} \ln \left(\frac{2}{\beta} \right) \right\rceil \quad \text{and} \quad t = 2 + \left\lceil \frac{9}{2\varepsilon} \ln \left(\frac{8}{\beta \cdot \delta} \right) \right\rceil$$

iii. For every $\beta \geq \beta_0$, $\text{StabilityHistogram}_{\mathcal{M}, b}$ has (a, β) -simultaneous accuracy for

$$a = 2 + \left\lceil \frac{9}{2\varepsilon} \ln \left(\frac{8n}{\beta \cdot \delta} \right) \right\rceil$$

iv. For $b \in [n]$, $\text{StabilityHistogram}_{\mathcal{M}, b}$ has running time

$$\tilde{O} \left(\frac{n}{\varepsilon} \cdot \log \frac{1}{\beta_0 \cdot \delta} \right) + O(n \log n \cdot \log |\mathcal{X}|)$$

\mathcal{M}	b	Running Time
GeometricMechanism	$1 + \lceil \frac{2}{\varepsilon} \ln \frac{2}{\delta} \rceil$	n/a
GeoSample	$1 + \lceil \frac{9}{2\varepsilon} \ln \frac{2}{\delta} \rceil$	$\tilde{O}(n^2 \cdot \log(1/\varepsilon)) + O(n \log n \cdot \log \mathcal{X})$
FastSample	$1 + \lceil \frac{9}{2\varepsilon} \ln \frac{4}{\delta} \rceil$	$\tilde{O}((n/\varepsilon) \cdot \log(1/(\beta_0 \cdot \delta))) + O(n \log n \cdot \log \mathcal{X})$

\mathcal{M}	(a, β) -Per-Query on $c_x(D) > t$		
	a	t	(a, β) -Simultaneous
GeometricMechanism	$\lceil \frac{2}{\varepsilon} \ln \frac{1}{\beta} \rceil$	$2 + \lceil \frac{2}{\varepsilon} \ln \frac{2}{\beta \cdot \delta} \rceil$	$2 + \lceil \frac{2}{\varepsilon} \ln \frac{2n}{\beta \cdot \delta} \rceil$
GeoSample	$\lceil \frac{9}{2\varepsilon} \ln \frac{1}{\beta} \rceil$	$2 + \lceil \frac{9}{2\varepsilon} \ln \frac{2}{\beta \cdot \delta} \rceil$	$2 + \lceil \frac{9}{2\varepsilon} \ln \frac{2n}{\beta \cdot \delta} \rceil$
FastSample	$\lceil \frac{9}{2\varepsilon} \ln \frac{2}{\beta} \rceil$	$2 + \lceil \frac{9}{2\varepsilon} \ln \frac{8}{\beta \cdot \delta} \rceil$	$2 + \lceil \frac{9}{2\varepsilon} \ln \frac{8n}{\beta \cdot \delta} \rceil$

Figure 4: The running time and errors of $\text{StabilityHistogram}_{\mathcal{M}, b}$ for the counting query algorithms of Section 4. Values shown are for a (ε, δ) -differentially private release where $\varepsilon, \delta \in \mathbb{N}^{-1}$. We assume $b \in [n]$ as otherwise the algorithm always outputs an empty histogram. For **FastSample**, the choice of parameters is specified in Theorem 5.6.

6 Improving the Running Time

While $\text{StabilityHistogram}$ has running time logarithmic in the universe size, it can only guarantee approximate differential privacy. In this section, we present an algorithm whose running time

depends only poly-logarithmically on the universe size while maintaining pure differential privacy based on the observation that most counts are 0 when $n \ll |\mathcal{X}|$; this is the same observation made by Cormode, Procopiuc, Srivastava and Tran [CPST12] to release private histograms that are sparse in expectation.

6.1 Sparse Histograms

We start by reducing the output length of $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ to release only the bins with the heaviest (or largest) counts (interpreted as a partial vector).

Algorithm 6.1. $\text{KeepHeavy}_{\mathcal{M}}(D)$ for $D \in \mathcal{X}^n$ where $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$

1. Let $\{(x, \tilde{c}_x)\}_{x \in \mathcal{X}} = \text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}(D)$.
2. Let x_1, \dots, x_{n+1} be the elements of \mathcal{X} with the largest counts in sorted order, i.e.

$$\tilde{c}_{x_1} \geq \tilde{c}_{x_2} \geq \dots \geq \tilde{c}_{x_{n+1}} \geq \max_{x \in \mathcal{X} \setminus \{x_1, \dots, x_{n+1}\}} \tilde{c}_x$$

3. Release $h = \{(x, \tilde{c}_x) : x \in \mathcal{X} \text{ and } \tilde{c}_x > \tilde{c}_{x_{n+1}}\} \in \mathcal{H}_{n,n}(\mathcal{X})$.⁴

Observe that the output length has been improved to $O(n \cdot (\log |\mathcal{X}| + \log n))$ bits compared to the $O(|\mathcal{X}| \cdot (\log |\mathcal{X}| + \log n))$ bits needed to represent the output of $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$.

Theorem 6.2. *Let $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be $(\varepsilon/2, 0)$ -differentially private for counting queries such that $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ has (a_1, β) -per-query accuracy and (a_2, β) -simultaneous accuracy with $a_1 \leq a_2$. Then $\text{KeepHeavy}_{\mathcal{M}} : \mathcal{X}^n \rightarrow \mathcal{H}_{n,n}(\mathcal{X})$ has the following properties:*

- i. $\text{KeepHeavy}_{\mathcal{M}}$ is $(\varepsilon, 0)$ -differentially private.*
- ii. $\text{KeepHeavy}_{\mathcal{M}}$ has $(a_1, 2\beta)$ -per-query accuracy on counts larger than $2a_2$.*
- iii. $\text{KeepHeavy}_{\mathcal{M}}$ has $(2a_2, \beta)$ -simultaneous accuracy.*

Unlike $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$, by taking $\mathcal{M} = \text{GeoSample}$, the algorithm $\text{KeepHeavy}_{\mathcal{M}}$ achieves $(O(\log(1/\beta)/\varepsilon), \beta)$ -per-query accuracy only on counts larger than $O(\log(|\mathcal{X}|/\beta)/\varepsilon)$. This loss is necessary for any algorithm that outputs a sparse histogram as we will show in Theorem 7.2.

Proof of i. Privacy follows from the $(\varepsilon, 0)$ -differential privacy of $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ (Theorem 5.2 Part i) along with differential privacy's closure under post-processing (Lemma 2.2). \square

To prove the remaining parts, we start with the following lemma.

Lemma 6.3. *Define the event $E = \{\forall x \in \mathcal{X} \quad |\tilde{c}_x - c_x(D)| \leq a_2\}$. Then $\Pr[E] \geq 1 - \beta$ and E implies that for all $x \in \mathcal{X}$ such that $c_x(D) > 2a_2$ we have $\tilde{c}_x > \tilde{c}_{x_{n+1}}$.*

⁴If instead \mathcal{M} had a continuous output distribution this last step is equivalent to releasing the n heaviest bins. However, in the discrete case, where ties can occur, from the set $A \cup \{q_0, \dots, q_n\}$ we cannot determine all bins with a count tied for the n -th heaviest as there may be many other noisy counts tied with \tilde{c}_{x_n} . As a result, we only output the bins with a strictly heavier count than $\tilde{c}_{x_{n+1}}$.

Proof. $\Pr[E] \geq 1 - \beta$ follows from the (a_2, β) -simultaneous accuracy of $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$.

Assume the event E . Then for all $x \in \mathcal{X}$ such that $c_x(D) > 2a_2$, we have $\tilde{c}_x > a_2$ and for all $x \in \mathcal{X}$ such that $c_x(D) = 0$, we have $\tilde{c}_x \leq a_2$. Because there are at most n distinct elements in D , we have $\tilde{c}_x > \tilde{c}_{x_{n+1}}$ for all $x \in \mathcal{X}$ such that $c_x(D) > 2a_2$. \square

Proof of Theorem 6.2 Part ii. Let $x \in \mathcal{X}$ such that $c_x(D) > 2a_2$. We have

$$\Pr(|h_x - c_x(D)| > a_1) \leq \Pr(\tilde{c}_x \leq \tilde{c}_{x_{n+1}}) + \Pr(|\tilde{c}_x - c_x(D)| > a_1) \leq 2\beta$$

by Lemma 6.3 and the (a_1, β) -per-query accuracy of $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$. \square

Proof of Theorem 6.2 Part iii. Assume the event E . By Lemma 6.3, for all $x \in \mathcal{X}$ such that $c_x(D) > 2a_2$ we have $\tilde{c}_x > \tilde{c}_{x_{n+1}}$ which implies $h_x = \tilde{c}_x$. Thus, $|h_x - c_x(D)| \leq a_2$. For the remaining $x \in \mathcal{X}$ we have $|h_x - c_x(D)| \leq 2a_2$ as $h_x = 0$ or $h_x = \tilde{c}_x$. \square

However, as described KeepHeavy still requires adding noise to the count of every bin. The following algorithm $\text{KH}'_{\mathcal{M}} : \mathcal{X}^n \rightarrow \mathcal{H}_{n,n}(\mathcal{X})$ simulates $\text{KeepHeavy}_{\mathcal{M}}$ by generating a candidate set of heavy bins from which only the heaviest are released. This candidate set is constructed from all bins with nonzero true count and a sample representing the bins with a true count of 0 that have the heaviest noisy counts.

Algorithm 6.4. $\text{KH}'_{\mathcal{M}}(D)$ for $D \in \mathcal{X}^n$ where $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ and $|\mathcal{X}| \geq 2n + 1^5$

1. Let $A = \{x \in \mathcal{X} : c_x(D) > 0\}$ and $m = |\mathcal{X} \setminus A|$.
2. Let $\{(x, \tilde{c}_x)\}_{x \in A} = \text{BasicHistogram}_{\mathcal{M}, A}(D)$.
3. Pick a uniformly random sequence (q_0, \dots, q_n) of distinct elements from $\mathcal{X} \setminus A$.
4. Sample $(\tilde{c}_{q_0}, \dots, \tilde{c}_{q_n})$ from the joint distribution of the order statistics $(Z_{(m)}, \dots, Z_{(m-n)})$ where Z_1, \dots, Z_m are i.i.d. $\mathcal{M}(0, U)$ random variables with $U \sim \text{Unif}([d]_+)$.
5. Sort the elements of $A \cup \{q_0, \dots, q_n\}$ as $x_1, \dots, x_{|A|+n+1}$ such that $\tilde{c}_{x_1} \geq \dots \geq \tilde{c}_{x_{|A|+n+1}}$.
6. Release $h = \{(x, \tilde{c}_x) : x \in \{x_1, \dots, x_n\} \text{ and } \tilde{c}_x > \tilde{c}_{x_{n+1}}\} \in \mathcal{H}_{n,n}(\mathcal{X})$.

Proposition 6.5. $\text{KH}'_{\mathcal{M}}(D)$ is identically distributed to $\text{KeepHeavy}_{\mathcal{M}}(D)$.

Proof. Let $\{\hat{c}_x\}_{x \in \mathcal{X}}$ be the noisy counts set by $\text{KeepHeavy}_{\mathcal{M}}(D)$ and let $\hat{x}_1, \dots, \hat{x}_{n+1}$ be the sorted ordering of the $n + 1$ heaviest bins by these counts. We have $\tilde{c}_x \sim \hat{c}_x$ for all $x \in A$ and the Z_i 's are identically distributed to $\{\hat{c}_x\}_{x \in \mathcal{X} \setminus A}$.

$\{(q_i, \tilde{c}_{q_i})\}_{i=0}^n$ is identically distributed to the $n + 1$ bins with heaviest counts of $\{(x, \hat{c}_x)\}_{x \in \mathcal{X} \setminus A}$ (breaking ties uniformly at random) due the noisy counts of the empty bins being independent and identically distributed. Therefore,

$$\begin{aligned} h &= \{(x, \tilde{c}_x) : x \in \{x_1, \dots, x_n\} \text{ and } \tilde{c}_x > \tilde{c}_{x_{n+1}}\} \\ &= \{(x, \tilde{c}_x) : x \in A \cup \{q_0, \dots, q_n\} \text{ and } \tilde{c}_x > \tilde{c}_{x_{n+1}}\} \\ &\sim \{(x, \hat{c}_x) : x \in \mathcal{X} \text{ and } \hat{c}_x > \hat{c}_{\hat{x}_{n+1}}\} \end{aligned}$$

which shows that $\text{KH}'_{\mathcal{M}}(D)$ is identically distributed to $\text{KeepHeavy}_{\mathcal{M}}(D)$. \square

⁵ $|\mathcal{X}| \geq 2n + 1$ ensures that $|\mathcal{X} \setminus A| \geq n + 1$. One can use $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}(D)$ when $|\mathcal{X}| \leq 2n$.

The original version of our paper [BV18] used an incorrect method for sampling the order statistics based on a claim that their conditional CDF is given by:

$$F_{Z_{(m-k)}|Z_{(m-k+1)}=v_{m-k+1}, \dots, Z_{(m)}=v_m}(z) = \begin{cases} (F(z)/F(v_{m-k+1}))^{m-k} & \text{if } z \leq v_{m-k+1} \\ 1 & \text{otherwise} \end{cases}$$

Unfortunately, this only holds for continuous distributions. Here we correct the error with a different method of sampling from the order statistics. In order to sample from the order statistics used by KH' we construct an algorithm that samples from binomial distributions to construct a histogram of $Z_{(m)}, \dots, Z_{(m-n)}$.

Proposition 6.6. *Let $n, d \in \mathbb{N}_+$ and $F : [n] \rightarrow [d]_+$ such that F is non-decreasing and $F(n) = d$. Let $m \in \mathbb{N}_+$ such that $m \geq n+1$. Let Z_1, \dots, Z_m be i.i.d. random variables over $[n]$ with cumulative distribution function $F(z)/d$ for all $z \in [n]$. Then the following algorithm $\text{OrdSample}_F(m)$ is identically distributed to the top $(n+1)$ order statistics $(Z_{(m)}, \dots, Z_{(m-n)})$.*

Algorithm 6.7. $\text{OrdSample}_F(m)$ for $m \in \mathbb{N}_+$ such that $m \geq n+1$ where $F : [n] \rightarrow [d]_+$ such that F is non-decreasing and $F(n) = d$

1. Let $L_{n+1} = 0$.
2. For v from n to 1 , do the following:
 - (a) Sample $\ell_v \sim \min\{\text{Bin}(m - L_{v+1}, p_v), n + 1 - L_{v+1}\}$ where $p_v = 1 - \frac{F(v-1)}{F(v)}$.
 - (b) Let $L_v = L_{v+1} + \ell_v$.
3. Let $\ell_0 = n + 1 - L_1$.
4. Return (c_0, \dots, c_n) such that the first ℓ_n values are n , the next ℓ_{n-1} values are $n-1$ and so on until the last ℓ_0 values are 0 .

Proof. For $v \in [n]$, let $Y_v = |\{i : Z_i = v\}|$. Using these random variables we can compute the order statistics $Z_{(m)}, \dots, Z_{(1)}$ by taking the first Y_n values to be n , the next Y_{n-1} to be $n-1$ and so on.

Since we only need the top $n+1$ order statistics, it suffices for us to stop calculating the Y_v 's once their sum exceeds $n+1$. That is, we can consider the random variables $Y'_v = \min\{Y_v, n+1 - \sum_{i=v+1}^n Y_i\}$. Notice that

$$Y'_v = \begin{cases} 0 & \text{if } v < v^* \\ n+1 - \sum_{i=v^*+1}^n Y_i & \text{if } v = v^* \\ Y_v & \text{if } v > v^* \end{cases} \quad \text{where } v^* = \max\{v \in [n] : \sum_{i=v}^n Y_i \geq n+1\}$$

Therefore, using (Y'_0, \dots, Y'_n) we can compute the order statistics $(Z_{(m)}, \dots, Z_{(m-n)})$ in the same manner as we did with the Y_v 's.

Observe that conditioned on $\{Y_n = y_n, Y_{n-1} = y_{n-1}, \dots, Y_{v+1} = y_{v+1}\}$, we have $Y_v \sim \text{Bin}(m - L_{v+1}, p_v)$ where $L_{v+1} = \sum_{i=v+1}^n Y_i$ is the number of "assigned" values of Z_j and

$$p_v = \Pr[Z_j = v | Z_j \leq v] = 1 - \Pr[Z_j \leq v-1 | Z_j \leq v] = 1 - \frac{F(v-1)}{F(v)}$$

We now prove by induction from $v = n$ down to $v = 1$ that

$$\ell_v \mid \ell_{v+1}, \dots, \ell_n \sim Y'_v \mid Y'_{v+1}, \dots, Y'_n$$

For the base case, observe that:

$$\ell_n \sim \min\{\text{Bin}(m, p_n), n + 1\} \sim Y'_n$$

For the induction step, we have that for $v \in [n - 1]_+$,

$$\begin{aligned} & \ell_v \mid \ell_{v+1}, \dots, \ell_n \\ & \sim \begin{cases} \min\{\text{Bin}(m - L_{v+1}, p_v), n + 1 - L_{v+1}\} & \text{if } L_{v+1} < n + 1 \\ 0 & \text{otherwise} \end{cases} \mid \ell_{v+1}, \dots, \ell_n \\ & \sim \begin{cases} \min\{\text{Bin}(m - \sum_{i=v+1}^n Y'_i, p_v), n + 1 - \sum_{i=v+1}^n Y'_i\} & \text{if } \sum_{i=v+1}^n Y'_i < n + 1 \\ 0 & \text{otherwise} \end{cases} \mid Y'_{v+1}, \dots, Y'_n \\ & \sim \begin{cases} \min\{\text{Bin}(m - \sum_{i=v+1}^n Y_i, p_v), n + 1 - \sum_{i=v+1}^n Y_i\} & \text{if } \sum_{i=v+1}^n Y_i < n + 1 \\ 0 & \text{otherwise} \end{cases} \mid Y'_{v+1}, \dots, Y'_n \\ & \sim \min\left\{Y_v, n + 1 - \sum_{i=v+1}^n Y'_i\right\} \mid Y'_{v+1}, \dots, Y'_n \\ & \sim Y'_v \mid Y'_{v+1}, \dots, Y'_n \end{aligned}$$

By construction, we have $\ell_0 \mid \ell_1, \dots, \ell_n \sim Y'_0 \mid Y'_1, \dots, Y'_n$. Therefore, $(\ell_0, \dots, \ell_n) \sim (Y'_0, \dots, Y'_n)$ and so this algorithm returns values (c_0, \dots, c_n) distributed as $(Z_{(m)}, \dots, Z_{(m-n)})$. \square

Unfortunately, the running time of this algorithm is $\Omega(m)$ as it requires $\Omega(m)$ bits to represent the probabilities in $\text{Bin}(m, p_v)$. So while $\text{KH}'_{\mathcal{M}}$ only has an output of length $O(n \cdot (\log |\mathcal{X}| + \log n))$, by using $\text{OrdSample}_{F_{\mathcal{M}}}$ to implement step 4, $\text{KH}'_{\mathcal{M}}$ has running time at least linear in $m \geq |\mathcal{X}| - n$. Indeed, this is necessary since the distribution of the order statistic $Z_{(m)}$ has probabilities that are exponentially small in m .⁶

6.2 An Efficient Approximation

To remedy the inefficiency of KH' we construct an efficient algorithm that approximates the output distribution of KH' by replacing OrdSample with an efficient algorithm whose output distribution is close in statistical distance to that of OrdSample , resulting in the following histogram algorithm:

Theorem 6.8. *Let $\varepsilon, \delta \in \mathbb{N}^{-1}$ and deterministic $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be $(\varepsilon/2, 0)$ -differentially private for counting queries. Let $|\mathcal{X}| \geq 2n + 1$. Then there exists an algorithm $\text{SparseHistogram}_{\mathcal{M}, \delta} : \mathcal{X}^n \rightarrow \mathcal{H}_{n,n}(\mathcal{X})$ with the following properties:*

- i. $\Delta(\text{KH}'_{\mathcal{M}}(D), \text{SparseHistogram}_{\mathcal{M}, \delta}(D)) \leq \delta$ for all $D \in \mathcal{X}^n$.
- ii. $\text{SparseHistogram}_{\mathcal{M}, \delta}$ is $(\varepsilon, (e^\varepsilon + 1) \cdot \delta)$ -differentially private.
- iii. The running time of $\text{SparseHistogram}_{\mathcal{M}, \beta_1, \delta}$ is

$$\tilde{O}(n \cdot \log |\mathcal{X}| \cdot (\log |\mathcal{X}| + \log(1/\delta))) + O(n) \cdot \left(\tilde{O}(\log d) + \text{Time}(\mathcal{M}) + \text{Time}(F_{\mathcal{M}})\right)$$

⁶Notice $\Pr[Z_{(m)} = 0] = \Pr[\mathcal{M}(0, U)]^m$. Thus, we need $\Omega(m)$ random bits to sample from $Z_{(m)}$.

Note that this algorithm only achieves $(\varepsilon, O(\delta))$ -differential privacy. By reducing δ , the algorithm better approximates KH' , at only the cost of increasing running time (polynomial in the bit length of δ). This is in contrast to most (ε, δ) -differentially private algorithms such as the stability based algorithm of Section 5.2, where one needs $n \geq \Omega(\log(1/\delta)/\varepsilon)$ to get any meaningful accuracy.

We will prove Theorem 6.8 in Section 6.3. Before that we will convert **SparseHistogram** to a pure differentially private algorithm by mixing it with random output, following Lemma 3.2.

Algorithm 6.9. $\text{PureSparseHistogram}_{\mathcal{M}, \varepsilon, \beta_1}(D)$ for $D \in \mathcal{X}^n$ where deterministic $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$, $\varepsilon, \beta_1 \in \mathbb{N}^{-1}$ s.t. $\beta_1 < 1$ and $|\mathcal{X}| \geq 2n + 1$

1. With probability $1 - \beta_1$ release $\text{SparseHistogram}_{\mathcal{M}, \delta}(D)$ with

$$\delta = \frac{\varepsilon}{3} \cdot \beta_1 \cdot \left(\frac{1}{3 \cdot |\mathcal{X}|} \right)^n$$

2. Otherwise

- (a) Draw (x_1, \dots, x_n) uniformly at random from \mathcal{X}^n .
- (b) Let Q be the set of distinct elements from (x_1, \dots, x_n) .
- (c) For each $q \in Q$, sample \tilde{c}_q uniformly at random from $[n]$.
- (d) Release $h = \{(q, \tilde{c}_q) : q \in Q \text{ and } \tilde{c}_q > 0\} \in \mathcal{H}_{n,n}(\mathcal{X})$.

Theorem 6.10. *Let $\varepsilon, \beta_1 \in \mathbb{N}^{-1}$ such that $\beta_1 < 1$ and $|\mathcal{X}| \geq 2n + 1$. Let deterministic $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be $(\varepsilon/2, 0)$ -differentially private for counting queries such that $\text{BasicHistogram}_{\mathcal{M}, \mathcal{X}}$ has (a_1, β_2) -per-query accuracy and (a_2, β_2) -simultaneous accuracy with $a_1 \leq a_2$.*

Then $\text{PureSparseHistogram}_{\mathcal{M}, \varepsilon, \beta_1} : \mathcal{X}^n \rightarrow \mathcal{H}_{n,n}(\mathcal{X})$ has the following properties:

- i. $\text{PureSparseHistogram}_{\mathcal{M}, \varepsilon, \beta_1}$ is $(\varepsilon, 0)$ -differentially private.*
- ii. $\text{PureSparseHistogram}_{\mathcal{M}, \varepsilon, \beta_1}$ has $(a_1, 2\beta_1 + 2\beta_2)$ -per-query accuracy on counts larger than $2a_2$.*
- iii. $\text{PureSparseHistogram}_{\mathcal{M}, \varepsilon, \beta_1}$ has $(2a_2, 2\beta_1 + \beta_2)$ -simultaneous accuracy.*
- iv. $\text{PureSparseHistogram}_{\mathcal{M}, \varepsilon, \beta_1}$ has running time*

$$\tilde{O}(n^2 \log^2 |\mathcal{X}| + n \cdot \log |\mathcal{X}| \cdot \log(1/(\beta_1 \varepsilon))) + O(n) \cdot (\tilde{O}(\log d) + \text{Time}(\mathcal{M}) + \text{Time}(F_{\mathcal{M}}))$$

Notice that the running time of this algorithm depends nearly quadratically on n . It is an open problem as to whether or not one can improve the nearly quadratic dependence in running time on n to nearly linear while maintaining the sparsity, privacy and accuracy guarantees achieved by this algorithm. The nearly quadratic running time dependence on n is due to approximating the distribution of each of the $O(n)$ order statistics in KH' to within a statistical distance that is exponentially small in n (in order to apply Lemma 3.2). See Section 6.3 for more details.

Proof of i. Define the distribution \mathcal{D} over $\mathcal{H}_{n,n}(\mathcal{X})$ as the histogram returned by second step of $\text{PureSparseHistogram}$. Notice that \mathcal{D} has full support over $\mathcal{H}_{n,n}(\mathcal{X})$ and is determined by our

sample from \mathcal{X}^n and $[n]^n$. Thus,

$$\min_{h' \in \mathcal{H}_{n,n}(\mathcal{X})} \Pr_{h \sim \mathcal{D}}[h = h'] \geq \frac{n!}{(n+1)^n \cdot |\mathcal{X}|^n} \geq \left(\frac{1}{3 \cdot |\mathcal{X}|} \right)^n$$

Privacy follows from Theorem 6.8 and Lemma 3.2 by taking $\mathcal{M} = \text{KH}'_{\mathcal{M}}$, $\mathcal{R} = \mathcal{H}_{n,n}(\mathcal{X})$, $\mathcal{M}' = \text{SparseHistogram}_{\mathcal{M},\delta}$, $\gamma = \beta_1$ and $\mathcal{D} = \mathcal{D}$ as

$$\delta \leq \frac{e^\varepsilon - 1}{e^\varepsilon + 1} \cdot \frac{\beta_1}{1 - \beta_1} \cdot \min_{h' \in \mathcal{H}_{n,n}(\mathcal{X})} \Pr_{h \sim \mathcal{D}}[h = h'] \quad \square$$

Proof of ii-iii. Let $D \in \mathcal{X}^n$. For any $x \in \mathcal{X}$ such that $c_x(D) > 2a_2$ define the set $G = \{h \in \mathcal{H}_{n,n}(\mathcal{X}) : |h_x - c_x(D)| \leq a_1\}$. By construction,

$$\Pr[\text{PureSparseHistogram}_{\mathcal{M},\varepsilon,\beta_1}(D) \in G] \geq \Pr[\text{SparseHistogram}_{\mathcal{M},\delta}(D) \in G] - \beta_1$$

where δ is defined in Algorithm 6.9. Notice that $\delta \leq \beta_1$. So by Theorem 6.8 Part *i*,

$$\Pr[\text{SparseHistogram}_{\mathcal{M},\delta}(D) \in G] - \beta_1 \geq \Pr[\text{KH}'_{\mathcal{M}}(D) \in G] - 2\beta_1$$

And by Theorem 6.2 Part *ii* and Proposition 6.5, we have

$$\Pr[\text{KH}'_{\mathcal{M}}(D) \in G] - 2\beta_1 \geq 1 - 2\beta_1 - 2\beta_2$$

Similarly, we can bound the simultaneous accuracy by using Theorem 6.2 Part *iii*. □

Proof of iv. The second step of `PureSparseHistogram` can be computed in $O(n \log n \cdot \log |\mathcal{X}|)$ time. The running time follows from Theorem 6.8 Part *iv* with

$$\log \frac{1}{\delta} = O\left(n \cdot \log |\mathcal{X}| + \log \frac{1}{\beta_1 \cdot \varepsilon}\right) \quad \square$$

Now we can use `PureSparseHistogram` with the counting query algorithms of Section 4.

Theorem 6.11. *Let $\varepsilon, \beta_0 \in \mathbb{N}^{-1}$ and $\mathcal{M} = \text{GeoSample}_{n,\varepsilon}$. Then*

- i. `PureSparseHistogram` $_{\mathcal{M},\varepsilon,\beta_0/4}$ is $(\varepsilon, 0)$ -differentially private.*
- ii. For every $\beta \geq \beta_0$, `PureSparseHistogram` $_{\mathcal{M},\varepsilon,\beta_0/4}$ has (a, β) -per-query accuracy on counts larger than t for*

$$a = \left\lceil \frac{9}{2\varepsilon} \ln \left(\frac{4}{\beta} \right) \right\rceil \quad \text{and} \quad t = 2 \cdot \left\lceil \frac{9}{2\varepsilon} \ln \left(\frac{4 \cdot |\mathcal{X}|}{\beta} \right) \right\rceil$$

- iii. For every $\beta \geq \beta_0$, `PureSparseHistogram` $_{\mathcal{M},\varepsilon,\beta_0/4}$ has (a, β) -simultaneous accuracy for*

$$a = 2 \cdot \left\lceil \frac{9}{2\varepsilon} \ln \left(\frac{2 \cdot |\mathcal{X}|}{\beta} \right) \right\rceil$$

- iv. `PureSparseHistogram` $_{\mathcal{M},\varepsilon,\beta_0/4}$ has running time*

$$\tilde{O}\left(n^2 \cdot \log^2 |\mathcal{X}| + n^2 \cdot \log(1/\varepsilon) + n \cdot \log |\mathcal{X}| \cdot \log(1/(\beta_0 \cdot \varepsilon))\right)$$

\mathcal{M}	Running Time		
GeoSample	$\tilde{O}(n^2 \cdot \log^2 \mathcal{X} + n^2 \cdot \log(1/\varepsilon) + n \cdot \log \mathcal{X} \cdot \log(1/(\beta \cdot \varepsilon)))$		
FastSample	$\tilde{O}(n^2 \cdot \log^2 \mathcal{X} + (n/\varepsilon) \cdot \log(\mathcal{X} /\beta) + n \cdot \log \mathcal{X} \cdot \log(1/(\beta \cdot \varepsilon)))$		
(a, β)-Per-Query on $c_x(D) > t$			
\mathcal{M}	a	t	(a, β)-Simultaneous
GeoSample	$\left\lceil \frac{9}{2\varepsilon} \ln \frac{4}{\beta} \right\rceil$	$2 \cdot \left\lceil \frac{9}{2\varepsilon} \ln \frac{4 \mathcal{X} }{\beta} \right\rceil$	$2 \cdot \left\lceil \frac{9}{2\varepsilon} \ln \frac{2 \mathcal{X} }{\beta} \right\rceil$
FastSample	$\left\lceil \frac{9}{2\varepsilon} \ln \frac{8}{\beta} \right\rceil$	$2 \cdot \left\lceil \frac{9}{2\varepsilon} \ln \frac{8 \mathcal{X} }{\beta} \right\rceil$	$2 \cdot \left\lceil \frac{9}{2\varepsilon} \ln \frac{4 \mathcal{X} }{\beta} \right\rceil$

Figure 5: The running time and errors of $\text{PureSparseHistogram}_{\mathcal{M}, \varepsilon, \beta/4}$ for the counting query algorithms of Section 4 where $\varepsilon, \beta \in \mathbb{N}^{-1}$. For per-query accuracy, the first value is the error a and the second value is the threshold t . Values shown are for a $(\varepsilon, 0)$ -differentially private release. For FastSample , we take $\gamma = \beta/(8|\mathcal{X}|)$.

6.3 Construction of SparseHistogram

We finish this section with the construction of SparseHistogram . Notice that if we implement KH' by using $\text{OrdSample}(m)$ to implement step 4, then OrdSample will have to sample from $\text{Bin}(i, p)$ for $i \geq |\mathcal{X}| - n$ and rational p . This in turn has probabilities as small as p^i , whose bit length is too large for us.

In order to keep our bit lengths manageable, we will only sample from a distribution close in statistical distance to the desired Binomial distribution. We will leverage the fact that a Binomial random variable can be represented as the sum of independent Bernoulli random variables. By representing the probability mass function of a Bernoulli as a vector, we can then compute the probability mass function of a Binomial by repeatedly convolving this vector.

Definition 6.12. Let $a, b \in \mathbb{R}^t$. We define the **convolution** $a * b \in \mathbb{R}^t$ such that $(a * b)_k = \sum_{i=0}^k a_i b_{k-i}$ for all $k \in [t-1]$. And we define **i -fold convolution of a** denoted $*^{(i)}(a)$ to be $i-1$ convolutions of a with itself.

Lemma 6.13. Let $Y_1, Y_2 \in [t]$ be independent discrete random variables with probability mass functions f_{Y_1} and f_{Y_2} represented as vectors in \mathbb{R}^{t+1} . Then $\Pr[Y_1 + Y_2 = k] = (f_{Y_1} * f_{Y_2})_k$ for all $k \in [t]$.

Corollary 6.14. Let $Z \sim \text{Bin}(m, p)$ and $Y \sim \text{Bern}(p)$ with probability mass functions f_Z and f_Y represented as vectors in \mathbb{R}^{m+1} . In particular, $f_Y = (1-p, p, 0, \dots, 0)$. Then we have $f_Z = *^{(m)}(f_Y)$.

The following algorithm will approximate i -fold convolution of a by using an algorithm similar to exponentiation by repeated squaring while truncating each intermediate result to keep its bit length manageable. The following lemma provides a bound on the error and on the running time.

Proposition 6.15. There is an algorithm $\text{ApproxConvExp}_{s,t}(a, i)$ such that for all $i, s, t \in \mathbb{N}_+$ and $a \in [s]^t$ such that $\|a\|_1 \leq s$:

- i. $\text{ApproxConvExp}_{s,t}(a, i) \in [s]^t$ and $\|\text{ApproxConvExp}_{s,t}(a, i)\|_1 \leq s$.

ii. The approximation does not change as t varies. In particular, let $d = \text{ApproxConvExp}_{s,t}(a, i)$. Then for $t' < t$

$$(d_0, \dots, d_{t'-1}) = \text{ApproxConvExp}_{s,t'}((a_0, \dots, a_{t'-1}), i)$$

iii. $\text{ApproxConvExp}_{s,t}(a, i)$ satisfies the accuracy bound

$$\left\| \frac{\text{ApproxConvExp}_{s,t}(a, i)}{s} - \frac{*^{(i)}(a)}{s^i} \right\|_1 \leq \frac{t}{s} \cdot (i - 1)$$

iv. $\text{ApproxConvExp}_{s,t}(a, i)$ has running time $O(\log i) \cdot \tilde{O}(t \cdot \log s) + O(\log^2 i)$.

Proof. The algorithm is defined as follows.

Algorithm 6.16. $\text{ApproxConvExp}_{s,t}(a, i)$ for $i, s, t \in \mathbb{N}_+$ and $a \in [s]^t$

1. If $i = 1$, stop and return a .
 2. Let $b = \text{ApproxConvExp}_{s,t}(a, \lfloor i/2 \rfloor)$.
 3. Compute $c = \begin{cases} (b * b)/s & \text{if } i \text{ is even} \\ (a * b * b)/s^2 & \text{if } i \text{ is odd} \end{cases}$
 4. Return $d = (\lfloor c_0 \rfloor, \lfloor c_1 \rfloor, \dots, \lfloor c_{t-1} \rfloor)$.
-

Proof of i. We prove this part by induction on i . The result is trivial for $i = 1$. Let $i > 1$. By the inductive hypothesis for $b = \text{ApproxConvExp}_{s,t}(a, \lfloor i/2 \rfloor)$, we have $\|b\|_1 \leq s$. If i is odd, then $c = (b * b)/s$ and we have $\|c\|_1 \leq \|b\|_1^2/s \leq s$. If i is even, then $c = (a * b * b)/s^2$ and we have $\|c\|_1 \leq \|a\|_1 \cdot \|b\|_1^2/s^2 \leq s$. Finally, for the output $d = \text{ApproxConvExp}_{s,t}(a, i) = (\lfloor c_0 \rfloor, \dots, \lfloor c_{t-1} \rfloor)$, we have $\|d\|_1 \leq \|c\|_1 \leq s$. \square

Proof of ii. Let $a' = (a_0, \dots, a_{t'-1})$. We prove this part by induction on i . The result is trivial for $i = 1$. Let $i > 1$. By the inductive hypothesis, for $b = \text{ApproxConvExp}_{s,t}(a, \lfloor i/2 \rfloor)$ and $b' = \text{ApproxConvExp}_{s,t'}(a', \lfloor i/2 \rfloor)$ we have $b' = (b_0, \dots, b_{t'-1})$. Then for $k \in [t' - 1]$ and i even

$$\begin{aligned} \text{ApproxConvExp}_{s,t}(a, i)_k &= \lfloor (b * b)_k / s \rfloor \\ &= \lfloor (b' * b')_k / s \rfloor = \text{ApproxConvExp}_{s,t'}(a', i)_k \end{aligned}$$

where the second equality holds because $(b * b)_k$ only depends on the first $k + 1$ terms of b (i.e. b_0, \dots, b_k) which are equal those of b' since $k < t'$. Similarly this induction holds for i odd. \square

Proof of iii. We prove this part by induction on i . The result is trivial for $i = 1$. Let $i > 1$. For i

odd, we can bound the error as

$$\begin{aligned}
\left\| \frac{d}{s} - \frac{*^{(i)}(a)}{s^i} \right\|_1 &\leq \left\| \frac{d-c}{s} \right\|_1 + \left\| \frac{c}{s} - \frac{*^{(i)}(a)}{s^i} \right\|_1 \\
&\leq \frac{t}{s} + \left\| \left(\frac{a}{s} * \frac{b}{s} * \frac{b}{s} \right) - \left(\frac{a}{s} * \frac{*^{(\lfloor i/2 \rfloor)}(a)}{s^{\lfloor i/2 \rfloor}} * \frac{*^{(\lfloor i/2 \rfloor)}(a)}{s^{\lfloor i/2 \rfloor}} \right) \right\|_1 \\
&\leq \frac{t}{s} + \left\| \frac{a}{s} \right\|_1 \cdot \left\| \frac{b}{s} + \frac{*^{(\lfloor i/2 \rfloor)}(a)}{s^{\lfloor i/2 \rfloor}} \right\|_1 \cdot \left\| \frac{b}{s} - \frac{*^{(\lfloor i/2 \rfloor)}(a)}{s^{\lfloor i/2 \rfloor}} \right\|_1 \\
&\leq \frac{t}{s} + 2 \cdot \left\| \frac{b}{s} - \frac{*^{(\lfloor i/2 \rfloor)}(a)}{s^{\lfloor i/2 \rfloor}} \right\|_1 \\
&\leq \frac{t}{s} \cdot (i-1)
\end{aligned}$$

with the last step by the inductive hypothesis. Similarly this induction holds for i even. \square

Proof of iv. The running time follows from the observation that a call to $\text{ApproxConvExp}_{s,t}(a, i)$ makes at most $O(\log i)$ recursive calls with each dominated by computing $\lfloor i/2 \rfloor$ and the convolution of up to three vectors in $[s]^t$ which can be done in time $\tilde{O}(t \cdot \log s)$ ([vzGG13] Corollary 8.27). \square

Using ApproxConvExp we can approximately sample from a Binomial distribution truncated to a specified upper bound.

Proposition 6.17. *There exists an algorithm $\text{ApproxBinSample}_{s,t}(m, p, q)$ such that for $m, q, s \in \mathbb{N}_+$ with $s \geq m$, $t \in \mathbb{N}$, $p \in [q]$ and $Z \sim \min\{\text{Bin}(m, p/q), t\}$, we have*

$$\Delta(\text{ApproxBinSample}_{s,t}(m, p, q), Z) \leq \frac{m \cdot (t+1) - t}{s}$$

In addition for every $\ell \in [t]$, every execution of $\text{ApproxBinSample}_{s,t}(m, p, q)$ that produces an output of ℓ has running time at most

$$\tilde{O}(\log q) + O(\log m) \cdot \tilde{O}(\ell \cdot \log s)$$

Let $a = (1 - p/q, p/q, 0, \dots, 0)$. By using $\text{ApproxConvExp}_{s,t}$, we can approximate $*^{(m)}(a)$ and therefore the CDF of Z for $k < t$. Then to approximately sample from Z we follow Lemma 2.13 by first generating a random uniform u and then outputting the smallest ℓ such that the CDF at ℓ is at least u . This would yield what is claimed in Proposition 6.17, except that the running time would depend nearly linearly on t instead of the specific output ℓ . To remedy this, we approximate the first t' terms of $*^{(m)}(a)$ for $t' = 1, 2, 4, \dots$ until we find an $\ell < t'$ such that CDF of Z at ℓ is at least u .

Proof. The algorithm is defined as follows.

Algorithm 6.18. $\text{ApproxBinSample}_{s,t}(m, p, q)$ for $m, s, q \in \mathbb{N}_+$ such that $s \geq m$, $t \in \mathbb{N}$ and $p \in [q]$

1. If $t = 0$, stop and return 0.
 2. Sample u uniformly at random from $[s]_+$.
 3. Let $t' = 1$ and $p' = \lfloor sp/q \rfloor$.
 4. While $t' < 2t$, do the following:
 - (a) Let $a \in [s]^{t'}$ such that $a_0 = s - p'$, $a_1 = p'$ and $a_k = 0$ for $k \geq 2$.
 - (b) Let $d^{(t')} = \text{ApproxConvExp}_{s,t'}(a, m)$.
 - (c) Let $F_0^{(t')} = d_0^{(t')}$ and $F_k^{(t')} = F_{k-1}^{(t')} + d_k^{(t')}$ for $k \in [t' - 1]_+$.
 - (d) For $\ell \in [\lfloor t'/2 \rfloor, \dots, t' - 1]$, do the following:
 - i. If $F_\ell^{(t')} \geq u$, stop and return $\min\{\ell, t\}$.
 - (e) Set t' to $2t'$.
 5. Return t .
-

Assume $t > 0$ (as accuracy is trivial otherwise). Notice that $\text{ApproxBinSample}_{s,t}(m, p) \in [t]$ and by Proposition 6.15 Parts *i-ii*, we have $F_k^{(t')} = F_k^{(t)}$ and $F_k^{(t)} \in [s]$ for $k < t' \leq t$. Then by construction for $\ell \in [t - 1]$,

$$\Pr[\text{ApproxBinSample}_{s,t}(m, p, q) = \ell] = \Pr\left[u \in \left(F_{\ell-1}^{(t)}, F_\ell^{(t)}\right]\right] = \frac{1}{s} \cdot d_\ell^{(t)}$$

Now, let $Z' \sim \min\{\text{Bin}(m, p'/s), t\}$. Let $Y \sim \text{Bern}(p/q)$ and $Y' \sim \text{Bern}(p'/s)$. Notice that

$$\Delta(Z, Z') \leq m \cdot \Delta(Y, Y') \leq \frac{m}{s}$$

and, by Corollary 6.14, $\Pr[Z' = k] = \binom{m}{k} (a)_k / s^m$ for $k \in [t - 1]$ where $a = (s - p', p', 0, \dots, 0) \in [s]^t$. Therefore, by Proposition 6.15 Part *iii*

$$\begin{aligned} \Delta(\text{ApproxBinSample}_{s,t}(m, p, q), Z) &\leq \Delta(Z, Z') + \Delta(\text{ApproxBinSample}_{s,t}(m, p, q), Z') \\ &\leq \frac{m}{s} + \left\| \frac{d^{(t)}}{s} - \frac{\binom{m}{a}}{s^m} \right\|_1 \\ &\leq \frac{m \cdot (t + 1) - t}{s} \end{aligned}$$

We consider each step to calculate the running time. Step 3 takes times $\tilde{O}(\log s + \log q)$. The i -th iteration of step 4 has $t' = 2^{i-1}$. Therefore, for the i -th iteration step 4b takes $O(\log m) \cdot \tilde{O}(2^i \cdot \log s)$ time and steps (4c-d) take $O(2^i \cdot \log s)$ time. Now, for the algorithm to output ℓ , it must halt in the $1 + \lceil \log_2(\ell + 1) \rceil$ -th iteration of step 4. Therefore, the overall running time is

$$\tilde{O}(\log s + \log q) + \sum_{i=1}^{1 + \lceil \log_2(\ell + 1) \rceil} O(\log m) \cdot \tilde{O}(2^i \cdot \log s) = \tilde{O}(\log q) + O(\log m) \cdot \tilde{O}(\ell \cdot \log s) \quad \square$$

Now we can modify `OrdSample` by replacing sampling from a binomial distribution with a call to `ApproxBinSample` to keep the bit lengths of its numbers from becoming too large, yielding an efficient algorithm whose output distribution is close to that of `OrdSample`.

Proposition 6.19. *Let $n, d \in \mathbb{N}_+$ and $F : [n] \rightarrow [d]_+$ such that F is non-decreasing and $F(n) = d$. Let $m \in \mathbb{N}_+$ such that $m \geq n + 1$ and $s \in \mathbb{N}_+$ such that $s \geq m$. Then the following algorithm `ApproxOrdSampleF(m, s)` satisfies*

$$\Delta(\text{OrdSample}_F(m), \text{ApproxOrdSample}_F(m, s)) \leq \frac{m \cdot (n^2 + 2n)}{s}$$

In addition, `ApproxOrdSampleF(m, s)` has running time

$$O(n) \cdot \left(\text{Time}(F) + \tilde{O}(\log d) \right) + O(\log m) \cdot \tilde{O}(n \cdot \log s)$$

where $\text{Time}(F)$ is the worst-case time to evaluate F .

The running time of `ApproxOrdSample` is exponentially faster than `OrdSample` as a function of m , which we take to be close to $|\mathcal{X}|$. However, `ApproxOrdSample` will still be the bottleneck of `PureSparseHistogram` as it has a nearly quadratic running time dependence on n (since we will take $s \geq 2^n$) as opposed to the other steps which only have a nearly linear dependence on n (see Theorem 6.10 Part *iv*).

Algorithm 6.20. `ApproxOrdSampleF(m, s)` for $m, s \in \mathbb{N}_+$ such that $m \geq n + 1$ and $s \geq m$ where $F : [n] \rightarrow [d]_+$ such that F is non-decreasing and $F(n) = d$

1. Let $L'_{n+1} = 0$.
2. For v from n to 1 , do the following:
 - (a) Let $\ell'_v = \text{ApproxBinSample}_{s, n+1-L'_{v+1}}(m - L'_{v+1}, F(v) - F(v-1), F(v))$.
 - (b) Let $L'_v = L'_{v+1} + \ell'_v$.
3. Let $\ell'_0 = n + 1 - L'_1$.
4. Return (c'_0, \dots, c'_n) such that the first ℓ'_n values are n , the next ℓ'_{n-1} values are $n - 1$ and so on until the last ℓ'_0 values are 0 .

Proof. Let (ℓ_0, \dots, ℓ_n) and (L_1, \dots, L_{n+1}) and be defined as in `OrdSampleF(m)`. And let $(\ell'_0, \dots, \ell'_n)$ and (L'_1, \dots, L'_{n+1}) be defined as in `ApproxOrdSampleF(m, s)`. Let $\mathcal{K}_v = \{(k_{v+1}, \dots, k_n) \in \mathbb{N}^{n-v} : \sum_{i=v+1}^n k_i \leq n + 1\}$. By Lemma 2.10 Part *iii* and induction,

$$\begin{aligned} \Delta(\text{ApproxOrdSample}_F(m, s), \text{OrdSample}_F(m)) &\leq \Delta((\ell_0, \dots, \ell_n), (\ell'_0, \dots, \ell'_n)) \\ &\leq \sum_{v=0}^n \max_{(k_{v+1}, \dots, k_n) \in \mathcal{K}_v} \Delta(\ell_v \mid \{\forall i > v \ \ell_i = k_i\}, \ell'_v \mid \{\forall i > v \ \ell'_i = k_i\}) \end{aligned}$$

Let $(k_{v+1}, \dots, k_n) \in \mathcal{K}_v$ and $K_{v+1} = \sum_{i=v+1}^n k_i$. Then for $v \in [n]_+$ we have

$$\ell'_v \mid \{\forall i > v \ \ell'_i = k_i\} \sim \text{ApproxBinSample}_{s, n+1-K_{v+1}}(m - K_{v+1}, F(v) - F(v-1), F(v))$$

and recall that

$$\ell_v \mid \{\forall i > v \ell_i = k_i\} \sim \min(\text{Bin}(m - K_{v+1}, p_v), n + 1 - K_{v+1}) \quad \text{where } p_v = \frac{F(v) - F(v-1)}{F(v)}$$

So by Proposition 6.17

$$\Delta(\ell_v \mid \{\forall i > v \ell_i = k_i\}, \ell'_v \mid \{\forall i > v \ell'_i = k_i\}) \leq \frac{m \cdot (n+2)}{s}$$

In addition, $\ell_0 \mid \{\forall i > 0 \ell_i = k_i\} \sim \ell'_0 \mid \{\forall i > 0 \ell'_i = k_i\}$. Therefore,

$$\Delta(\text{ApproxOrdSample}_F(m, s), \text{OrdSample}_F(m)) \leq \frac{m \cdot (n^2 + 2n)}{s}$$

The running time of this algorithm is dominated by the n calls to **ApproxBinSample** and the evaluation of F on all $n+1$ points. The other steps are simple arithmetic on $O(\log n)$ bit numbers and the final step can be done in $O(n \log n)$ time.

Because $\sum_{v=1}^n \ell'_v \leq n+1$, the running time to all calls of **ApproxBinSample** takes time

$$O(n) \cdot \tilde{O}(\log d) + \sum_{v=1}^n O(\log m) \cdot \tilde{O}(\ell'_v \cdot \log s) \leq O(n) \cdot \tilde{O}(\log d) + O(\log m) \cdot \tilde{O}(n \cdot \log s) \quad \square$$

We are ready to state the algorithm **SparseHistogram** and show it satisfies Theorem 6.8. It is identical to KH' except we replace sampling of order statistics with a call to **ApproxOrdSample**.

Algorithm 6.21. $\text{SparseHistogram}_{\mathcal{M}, \delta}(D)$ for $D \in \mathcal{X}^n$ where deterministic $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$, $\delta \in \mathbb{N}^{-1}$ and $|\mathcal{X}| \geq 2n+1$

1. Let $A = \{x \in \mathcal{X} : c_x(D) > 0\}$ and $m = |\mathcal{X} \setminus A|$.
2. Let $\{(x, \tilde{c}_x)\}_{x \in A} = \text{BasicHistogram}_{\mathcal{M}, A}(D)$.
3. Pick a uniformly random sequence (q_0, \dots, q_n) of distinct elements from $\mathcal{X} \setminus A$.
4. Let $(\tilde{c}_{q_0}, \dots, \tilde{c}_{q_n}) = \text{ApproxOrdSample}_{F_{\mathcal{M}}}(m, s)$ where $s = (n^2 + 2n) \cdot |\mathcal{X}| / \delta$.
5. Sort the elements of $A \cup \{q_0, \dots, q_n\}$ as $x_1, \dots, x_{|A|+n+1}$ such that $\tilde{c}_{x_1} \geq \dots \geq \tilde{c}_{x_{|A|+n+1}}$.
6. Release $h = \{(x, \tilde{c}_x) : x \in \{x_1, \dots, x_n\} \text{ and } \tilde{c}_x > \tilde{c}_{x_{n+1}}\} \in \mathcal{H}_{n,n}(\mathcal{X})$.

Theorem 6.8 (restated). *Let deterministic $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be $(\varepsilon/2, 0)$ -differentially private for counting queries. And let $\delta \in \mathbb{N}^{-1}$ and $|\mathcal{X}| \geq 2n+1$. Then $\text{SparseHistogram}_{\mathcal{M}, \beta_1, \delta} : \mathcal{X}^n \rightarrow \mathcal{H}_{n,n}(\mathcal{X})$ has the following properties:*

- i. $\Delta(\text{KH}'_{\mathcal{M}}(D), \text{SparseHistogram}_{\mathcal{M}, \delta}(D)) \leq \delta$ for all $D \in \mathcal{X}^n$.
- ii. $\text{SparseHistogram}_{\mathcal{M}, \delta}$ is $(\varepsilon, (e^\varepsilon + 1) \cdot \delta)$ -differentially private.
- iii. The running time of $\text{SparseHistogram}_{\mathcal{M}, \delta}$ is

$$\tilde{O}(n \cdot \log |\mathcal{X}| \cdot (\log |\mathcal{X}| + \log(1/\delta))) + O(n) \cdot (\tilde{O}(\log d) + \text{Time}(\mathcal{M}) + \text{Time}(F_{\mathcal{M}}))$$

Proof of i. Let $\text{KH}^* : \mathcal{X}^n \rightarrow \mathcal{H}_{n,2n+1}(\mathcal{X})$ be the algorithm KH' except, (if it passes the first step) instead of releasing the heaviest bins, KH^* releases the bins for all elements of $A \cup \{q_0, \dots, q_n\}$ (i.e. KH^* releases (x, \tilde{c}_x) for all $x \in A$ and (q_i, \tilde{c}_{q_i}) for all $i \in [n]$). Similarly, we define SparseHistogram^* with respect to SparseHistogram .

Notice that KH^* and SparseHistogram^* have the same distribution over the bins with nonzero true count. Only on the bins with counts sampled using OrdSample and ApproxOrdSample respectively do their output distributions differ. As a result, we can apply Proposition 6.19 to the output distributions of KH^* and SparseHistogram^* . So for all $D \in \mathcal{X}^n$

$$\Delta(\text{KH}_{\mathcal{M}}^*(D), \text{SparseHistogram}_{\mathcal{M},\delta}^*(D)) \leq (n^2 + 2n) \cdot \frac{|\mathcal{X}|}{s} = \delta$$

Now we consider the effect of keeping the heaviest counts. Define $T : \mathcal{H}_{n,2n+1}(\mathcal{X}) \rightarrow \mathcal{H}_{n,n}(\mathcal{X})$ to be the function that sets counts not strictly larger than the $(n+1)$ -heaviest count of its input to 0. Notice that $T \circ \text{KH}^* \sim \text{KH}'$ and $T \circ \text{SparseHistogram}^* \sim \text{SparseHistogram}$. So for all $D \in \mathcal{X}^n$, by Lemma 2.10 Part 2,

$$\begin{aligned} \Delta(\text{KH}'_{\mathcal{M}}(D), \text{SparseHistogram}_{\mathcal{M},\delta}(D)) &= \Delta(T(\text{KH}_{\mathcal{M}}^*(D)), T(\text{SparseHistogram}_{\mathcal{M},\delta}^*(D))) \\ &\leq \Delta(\text{KH}_{\mathcal{M}}^*(D), \text{SparseHistogram}_{\mathcal{M},\delta}^*(D)) \\ &\leq \delta \end{aligned} \quad \square$$

Proof of ii. Let D and D' be neighboring datasets. Let $S \subseteq \mathcal{H}_{n,n}(\mathcal{X})$. By the previous part, Theorem 6.2 Part *i* and Proposition 6.5,

$$\begin{aligned} \Pr[\text{SparseHistogram}_{\mathcal{M},\delta}(D) \in S] &\leq \Pr[\text{KH}'_{\mathcal{M}}(D) \in S] + \delta \\ &\leq e^\varepsilon \cdot \Pr[\text{KH}'_{\mathcal{M}}(D') \in S] + \delta \\ &\leq e^\varepsilon \cdot (\Pr[\text{SparseHistogram}_{\mathcal{M},\delta}(D) \in S] + \delta) + \delta \end{aligned}$$

Therefore, $\text{SparseHistogram}_{\mathcal{M},\delta}$ is $(\varepsilon, (e^\varepsilon + 1) \cdot \delta)$ -differentially private. \square

Proof of iii. We consider the running time at each step. Steps 1, 5 and 6 take time $O(n \log n \cdot \log |\mathcal{X}|)$. Step 3 can be done in time $O(n \log n \cdot \log^2 |\mathcal{X}|)$ (see Appendix A). For step 2, by Theorem 5.2 Part *i*, the call to $\text{BasicHistogram}_{\mathcal{M},A}(D)$ takes time

$$O(n \log n \cdot \log |\mathcal{X}| + n \cdot \log d + n \cdot \text{Time}(\mathcal{M}))$$

Notice that s can be computed in $\tilde{O}(\log n + \log |\mathcal{X}| + \log(1/\delta))$ time and has bit length $O(\log n + \log |\mathcal{X}| + \log(1/\delta))$. Thus, by Proposition 6.19 the call to $\text{ApproxOrdSample}_{F_{\mathcal{M}}}(m, s)$ in step 4 can be computed in time

$$\tilde{O}(n \cdot \log |\mathcal{X}| \cdot (\log |\mathcal{X}| + \log(1/\delta))) + O(n) \cdot (\tilde{O}(\log d) + \text{Time}(F_{\mathcal{M}}))$$

Therefore, overall $\text{SparseHistogram}_{\mathcal{M},\delta}$ has the desired running time. \square

7 Lower Bounds

In this section, we prove a lower bound on the per-query accuracy of histogram algorithms whose outputs are restricted to $\mathcal{H}_{\infty,n'}(\mathcal{X})$ (i.e. sparse histograms) using a packing argument [HT10, BBKN14]. First, for completeness we state and reprove existing lower bounds for per-query accuracy and simultaneous accuracy as well as generalize them to the case of $\delta > 0$.

Theorem 7.1 (following [HT10, BBKN14]). Let $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{H}_{\infty, |\mathcal{X}|}(\mathcal{X})$ be (ε, δ) -differentially private and $\beta \in (0, 1/2]$.

i. If \mathcal{M} has (a, β) -per-query accuracy, then

$$a \geq \frac{1}{2} \cdot \min \left\{ \frac{1}{\varepsilon} \ln \left(\frac{1}{4\beta} \right) - 1, \frac{1}{\varepsilon} \ln \left(\frac{\varepsilon}{4\delta} \right) - 1, n \right\}$$

ii. If \mathcal{M} has (a, β) -simultaneous accuracy, then

$$a \geq \frac{1}{2} \cdot \min \left\{ \frac{1}{\varepsilon} \ln \left(\frac{|\mathcal{X}| - 1}{4\beta} \right) - 1, \frac{1}{\varepsilon} \ln \left(\frac{\varepsilon}{4\delta} \right) - 1, n \right\}$$

Proof of i. Assume $a < n/2$. Let $x, x_0 \in \mathcal{X}$ such that $x \neq x_0$. Define the dataset $D' \in \mathcal{X}^n$ such that all rows are x_0 . And define the dataset D such that the first $m = \lfloor 2a \rfloor + 1$ rows are x and the remaining $n - m$ rows are x_0 . Notice that $\Pr[|\mathcal{M}(D)_x - c_x(D)| > a] \leq \beta$ by the (a, β) -per-query accuracy of \mathcal{M} . By Lemma 2.3 and the fact that $c_x(D) > 2a$ while $c_x(D') = 0$,

$$\begin{aligned} \Pr[|\mathcal{M}(D)_x - c_x(D)| > a] &\geq e^{-m\varepsilon} \cdot \Pr[|\mathcal{M}(D')_x - c_x(D)| > a] - \delta/\varepsilon \\ &\geq e^{-m\varepsilon} \cdot \Pr[|\mathcal{M}(D')_x - c_x(D')| \leq a] - \delta/\varepsilon \\ &\geq e^{-m\varepsilon} \cdot (1 - \beta) - \delta/\varepsilon \end{aligned}$$

Therefore,

$$e^{-(2a+1)\varepsilon} \leq \frac{1}{1-\beta} \cdot \left(\beta + \frac{\delta}{\varepsilon} \right) \leq 4 \cdot \max \left\{ \beta, \frac{\delta}{\varepsilon} \right\} \quad \square$$

Proof of ii. Assume $a < n/2$. Let $x_0 \in \mathcal{X}$. For each $x \in \mathcal{X}$ define the dataset $D^{(x)} \in \mathcal{X}^n$ such that the first $m = \lfloor 2a \rfloor + 1$ rows are x and the remaining $n - m$ rows are x_0 . For all $x \in \mathcal{X}$, let

$$G_x = \{h \in \mathcal{H}_{\infty, |\mathcal{X}|}(\mathcal{X}) : \forall x' \in \mathcal{X} \quad |h_{x'} - c_{x'}(D^{(x)})| \leq a\}$$

By Lemma 2.3, for all $x \in \mathcal{X}$

$$\begin{aligned} \Pr[\mathcal{M}(D^{(x_0)}) \in G_x] &\geq e^{-m\varepsilon} \cdot \Pr[\mathcal{M}(D^{(x)}) \in G_x] - \delta/\varepsilon \\ &\geq e^{-m\varepsilon} \cdot (1 - \beta) - \delta/\varepsilon \end{aligned}$$

Notice that $\Pr[\mathcal{M}(D^{(x_0)}) \notin G_{x_0}] \leq \beta$ and $\{G_x\}_{x \in \mathcal{X}}$ is a collection of disjoint sets. Then

$$\begin{aligned} \Pr[\mathcal{M}(D^{(x_0)}) \notin G_{x_0}] &\geq \sum_{x \in \mathcal{X}: x \neq x_0} \Pr[\mathcal{M}(D^{(x_0)}) \in G_x] \\ &\geq (|\mathcal{X}| - 1) \cdot (e^{-m\varepsilon} \cdot (1 - \beta) - \delta/\varepsilon) \end{aligned}$$

Therefore,

$$e^{-(2a+1)\varepsilon} \leq \frac{1}{1-\beta} \cdot \left(\frac{\beta}{|\mathcal{X}| - 1} + \frac{\delta}{\varepsilon} \right)$$

which implies the desired lower bound. □

We now state and prove our lower bound for privately releasing sparse histograms.

Theorem 7.2. Let $\mathcal{M} : \mathcal{X}^n \rightarrow \mathcal{H}_{\infty, n'}(\mathcal{X})$ be (ε, δ) -differentially private with (a, β) -per-query accuracy with $\beta \in (0, 1/2]$. Then

$$a \geq \frac{1}{2} \cdot \min \left\{ \frac{1}{2\varepsilon} \ln \left(\frac{|\mathcal{X}|}{16\beta n'} \right) - 1, \frac{1}{\varepsilon} \ln \left(\frac{\varepsilon}{4\delta} \right) - 1, n \right\}$$

The histogram algorithms of Sections 5.2 and 6 achieve $(O(\log(1/\beta)/\varepsilon), \beta)$ -per-query accuracy on large enough counts. However, on smaller counts we can only guarantee (a, β) -per-query accuracy with $a = O(\log(1/(\beta\delta))/\varepsilon)$ and $a = O(\log(|\mathcal{X}|/\beta)/\varepsilon)$ for algorithms from Sections 5.2 and 6 respectively (taking threshold $b = O(\log(1/\delta)/\varepsilon)$ in Section 5.2). Theorem 7.2 shows these bounds are the best possible, up to constant factors, when $|\mathcal{X}| \geq (n')^2$, $\varepsilon^2 \geq \delta$ and $\beta \geq \delta$.

Proof. Assume $a < n/2$. Let $x_0 \in \mathcal{X}$. For each $x \in \mathcal{X}$ define the dataset $D^{(x)} \in \mathcal{X}^n$ such that the first $m = \lceil 2a \rceil$ rows are x and the remaining $n - m$ rows are x_0 . By definition of (a, β) -per-query accuracy and the fact that $c_x(D^{(x)}) \geq 2a$, we have

$$\Pr \left[\mathcal{M}(D^{(x)})_x \geq a \right] \geq \Pr \left[\left| \mathcal{M}(D^{(x)})_x - c_x(D^{(x)}) \right| \leq a \right] \geq 1 - \beta$$

Then, by Lemma 2.3 and that $D^{(x)}$ is at distance at most m from $D^{(x_0)}$, we have

$$\Pr \left[\mathcal{M}(D^{(x_0)})_x \geq a \right] \geq (1 - \beta)e^{-m\varepsilon} - \delta/\varepsilon$$

Thus, by linearity of expectations

$$\mathbb{E} \left[\left| \left\{ x \in \mathcal{X} : \mathcal{M}(D^{(x_0)})_x \geq a \right\} \right| \right] \geq |\mathcal{X}| \cdot ((1 - \beta)e^{-m\varepsilon} - \delta/\varepsilon)$$

On the other hand, as $\mathcal{M}(D^{(x_0)}) \in \mathcal{H}_{\infty, n'}(\mathcal{X})$ we have

$$\mathbb{E} \left[\left| \left\{ x \in \mathcal{X} : \mathcal{M}(D^{(x_0)})_x \geq a \right\} \right| \right] \leq n'$$

Therefore,

$$e^{-\lceil 2a \rceil \cdot \varepsilon} \leq \frac{1}{1 - \beta} \cdot \left(\frac{n'}{|\mathcal{X}|} + \frac{\delta}{\varepsilon} \right)$$

which along with $\lceil 2a \rceil \leq 2a + 1$ implies the lower bound of

$$a \geq \frac{1}{2} \cdot \min \left\{ \frac{1}{\varepsilon} \ln \left(\frac{|\mathcal{X}|}{4n'} \right) - 1, \frac{1}{\varepsilon} \ln \left(\frac{\varepsilon}{4\delta} \right) - 1, n \right\}$$

Therefore, along with Theorem 7.1 Part *i*, we have

$$\begin{aligned} a &\geq \frac{1}{2} \cdot \min \left\{ \max \left\{ \frac{1}{\varepsilon} \ln \left(\frac{|\mathcal{X}|}{4n'} \right) - 1, \frac{1}{\varepsilon} \ln \left(\frac{1}{4\beta} \right) - 1 \right\}, \frac{1}{\varepsilon} \ln \left(\frac{\varepsilon}{4\delta} \right) - 1, n \right\} \\ &\geq \frac{1}{2} \cdot \min \left\{ \frac{1}{2\varepsilon} \ln \left(\frac{|\mathcal{X}|}{16\beta n'} \right) - 1, \frac{1}{\varepsilon} \ln \left(\frac{\varepsilon}{4\delta} \right) - 1, n \right\} \quad \square \end{aligned}$$

8 Better Per-Query Accuracy via Compact, Non-Sparse Representations

In this section, we present a histogram algorithm whose running time is poly-logarithmic in $|\mathcal{X}|$, but, unlike Algorithm 6.9, is able to achieve $(O(\log(1/\beta)/\varepsilon), \beta)$ -per query accuracy. It will output a histogram from a properly chosen family of succinctly representable histograms. This family necessarily contains histograms that have many nonzero counts to avoid the lower bound of Theorem 7.2.

8.1 The Family of Histograms

We start by defining this family of histograms.

Lemma 8.1. *Let deterministic $\mathcal{M}_0 : [d_0]_+ \rightarrow [n]$, $d_0 = 2^{2 \cdot 3^\ell}$ for some $\ell \in \mathbb{N}$, $d_0 \geq |\mathcal{X}|$ and $U \sim \text{Unif}([d_0]_+)$. There exists a multiset of histograms $\mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$ satisfying:*

- i. Let $g \sim \text{Unif}(\mathcal{G}_{\mathcal{M}_0}(\mathcal{X}))$. For all $x \in \mathcal{X}$, the marginal distribution g_x is distributed according to $\mathcal{M}_0(U)$.*
- ii. Let $g \sim \text{Unif}(\mathcal{G}_{\mathcal{M}_0}(\mathcal{X}))$. For all $B \subseteq \mathcal{X}$ such that $|B| \leq n + 1$ and for all $c \in [n]^B$*

$$\Pr[\forall x \in B \quad g_x = c_x] = \prod_{x \in B} \Pr[g_x = c_x]$$

- iii. For all $g \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$, the histogram g can be represented by a string of length $O(n \cdot \log d_0)$ and given this representation for all $x \in \mathcal{X}$ the count g_x can be evaluated in time*

$$O(n) \cdot \tilde{O}(\log d_0) + \text{Time}(\mathcal{M}_0)$$

- iv. For all $A \subseteq \mathcal{X}$ such that $|A| \leq n$ and $c \in [n]^A$ sampling a histogram h uniformly at random from $\{g \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X}) : \forall x \in A \quad g_x = c_x\}$ can be done in time*

$$O(n) \cdot \text{Time}(\mathcal{S}) + \tilde{O}(n \cdot \log d_0)$$

where $\text{Time}(\mathcal{S})$ is the maximum time over $v \in [n]$ to sample from the distribution $\mathcal{S}_v \sim \text{Unif}(\{u_0 \in [d_0]_+ : \mathcal{M}_0(u_0) = v\})$.

Proof. (Construction) Let $\mathcal{G}'_{\mathcal{M}_0}$ be the set of all degree at most n polynomials over the finite field \mathbb{F}_{d_0} . Now, $\mathcal{G}'_{\mathcal{M}_0}$ is a $(n + 1)$ -wise independent hash family mapping \mathbb{F}_{d_0} to \mathbb{F}_{d_0} . That is, $p \sim \text{Unif}(\mathcal{G}'_{\mathcal{M}_0})$ has the following properties:

- Let $x \in \mathbb{F}_{d_0}$. Then $p(x) \sim \text{Unif}(\mathbb{F}_{d_0})$.
- Let $x_0, \dots, x_n \in \mathbb{F}_{d_0}$ be distinct. Then the random variables $p(x_0), \dots, p(x_n)$ are independent.

And given any function $p_g \in \mathcal{G}'_{\mathcal{M}_0}$ we construct a histogram $g \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$ by using $p_g(x)$ as the randomness for \mathcal{M}_0 . More specifically, let $T : \mathbb{F}_{d_0} \rightarrow [d_0]_+$ be a bijection and for all $x \in \mathcal{X}$, define

$$g_x = \mathcal{M}_0(T(p_g(x)))$$

By construction, $g \sim \text{Unif}(\mathcal{G}_{\mathcal{M}_0}(\mathcal{X}))$ if and only if $p_g \sim \text{Unif}(\mathcal{G}'_{\mathcal{M}_0})$.

Proof of i. Let $g \sim \text{Unif}(\mathcal{G}_{\mathcal{M}_0}(\mathcal{X}))$ and $U \sim \text{Unif}([d_0]_+)$. Then $p_g \sim \text{Unif}(\mathcal{G}'_{\mathcal{M}_0})$ and because $\mathcal{G}'_{\mathcal{M}_0}$ is a $(n + 1)$ -wise independent hash family, for all $x \in \mathcal{X} \subseteq [d_0]_+$, $T(p_g(x)) \sim U$ which implies $g_x = \mathcal{M}_0(T(p_g(x))) \sim \mathcal{M}_0(U)$. \square

Proof of ii. Let $g \sim \text{Unif}(\mathcal{G}_{\mathcal{M}_0}(\mathcal{X}))$. Because $\mathcal{G}'_{\mathcal{M}_0}$ is a $(n + 1)$ -wise independent hash family, for all $B \subseteq \mathcal{X}$ such that $|B| \leq n + 1$ and for all $c \in [n]^B$

$$\begin{aligned} \Pr[\forall x \in B \quad g_x = c_x] &= \Pr[\forall x \in B \quad \mathcal{M}_0(T(p_g(x))) = c_x] \\ &= \prod_{x \in B} \Pr[\mathcal{M}_0(T(p_g(x))) = c_x] \\ &= \prod_{x \in B} \Pr[g_x = c_x] \end{aligned} \quad \square$$

Proof of iii. By choice of d_0 , we have $\mathbb{F}_{d_0} \simeq \mathbb{F}_2[x]/(x^{2 \cdot 3^\ell} + x^{3^\ell} + 1)$ [Lin99]. Thus, elements of \mathbb{F}_{d_0} can be represented by a polynomial of degree at most $\log d_0 - 1$ over \mathbb{F}_2 which requires $\log d_0$ bits. Arithmetic operations (addition, multiplication and inverse) of elements in \mathbb{F}_{d_0} can be done in time $\tilde{O}(\log d_0)$ ([vzGG13] Corollary 11.11). Also, this encoding defines an efficient bijection T between \mathbb{F}_{d_0} and $[d_0]_+$ by also interpreting the string as the binary representation of an element in $[d_0]_+$ offset by 1.

For all $g \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$, g can be represented by the coefficients of p_g . This representation can be encoded in $O(n \cdot \log d_0)$ bits. And given this encoding, the time to compute $g_x = \mathcal{M}_0(T(p_g(x)))$ follows from its construction. \square

Proof of iv. Let $A \subseteq \mathcal{X}$ such that $|A| \leq n$ and $c \in [n]^A$. Given A and c , we can sample $h \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$ given by the coefficients $a_0, \dots, a_n \in \mathbb{F}_{d_0}$ uniformly at random from $\{g \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X}) : \forall x \in A \ g_x = c_x\}$ with the following steps:

1. For each $x \in A$, sample u_x from the distribution \mathcal{S}_x .
2. Let $B \subseteq \mathcal{X}$ such that $A \subseteq B$ and $|B| = n + 1$. For all $x \in B \setminus A$, sample u_x uniformly at random from \mathbb{F}_{d_0} .
3. Take the coefficients $a_0, \dots, a_n \in \mathbb{F}_{d_0}$ to be the coefficients of the interpolating polynomial over \mathbb{F}_{d_0} given the set of points (x, u_x) for all $x \in B$.

We first prove correctness. Notice this procedure can only return a histogram $h \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$ such that $h_x = c_x$ for all $x \in A$ as the interpolating polynomial always exists. Now, let h be any such histogram. Then

$$\begin{aligned} \Pr[\text{Sampling } h] &= \Pr[(a_0, \dots, a_n) \text{ are the coefficients of } p_h] \\ &= \Pr[\forall x \in B \ u_x = p_h(x)] \\ &= \prod_{x \in B} \Pr[u_x = p_h(x)] \\ &= \left(\prod_{x \in A} \frac{1}{|\text{supp}(\mathcal{S}_x)|} \right) \cdot \left(\frac{1}{d_0} \right)^{|B \setminus A|} \end{aligned}$$

Therefore, these steps output $h \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$ uniformly at random such that $h_x = c_x$ for all $x \in A$.

Because $|A| \leq n$, the first step takes time $O(n) \cdot \text{Time}(\mathcal{S})$. Because $|B| = n + 1$, the second step takes time $O(n \log n \cdot \log d_0)$. Polynomial interpolation of $n + 1$ points over \mathbb{F}_{d_0} takes time $\tilde{O}(n \cdot \log d_0)$ ([vzGG13] Corollary 10.12). Thus, we have the desired running time overall. \square

8.2 The Algorithm

For our algorithm to have the correct marginal distributions over all bins we first compute the noisy counts for the nonzero bins using an algorithm \mathcal{M} that is differentially private for counting queries and then randomly pick a histogram from our family that is consistent with these computed counts. However, for technical reasons (e.g. requiring $d_0 \geq |\mathcal{X}|$) we allow our family to be defined in terms of an algorithm \mathcal{M}_0 that approximates \mathcal{M} ; we refer to \mathcal{M}_0 as the *empty-bin sampler*.

Algorithm 8.2. $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}(D)$ for $D \in \mathcal{X}^n$ where deterministic $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ and deterministic $\mathcal{M}_0 : [d_0]_+ \rightarrow [n]$ such that $d_0 = 2^{2 \cdot 3^\ell}$ for some $\ell \in \mathbb{N}$ and $d_0 \geq |\mathcal{X}|$

1. Let $A = \{x \in \mathcal{X} : c_x(D) > 0\}$.
 2. Let $\{(x, \tilde{c}_x)\}_{x \in A} = \text{BasicHistogram}_{\mathcal{M}, A}(D)$.
 3. Release h drawn uniformly at random from $\{g \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X}) : \forall x \in A \ g_x = \tilde{c}_x\}$.
-

Theorem 8.3. *Let deterministic $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ be $(\varepsilon_1/2, 0)$ -differentially private for counting queries and have (a, β) -accuracy. Let $\mathcal{M}_0 : [d_0]_+ \rightarrow [n]$ be deterministic, $d_0 = 2^{2 \cdot 3^\ell}$ for some $\ell \in \mathbb{N}$ and $d_0 \geq |\mathcal{X}|$. Assume $\Pr[\mathcal{M}_0(U_0) \leq a] \geq 1 - \beta$ and for all $c \in [n]$*

$$e^{-\varepsilon_2} \cdot \Pr[\mathcal{M}_0(U_0) = c] \leq \Pr[\mathcal{M}(0, U) = c] \leq e^{\varepsilon_3} \cdot \Pr[\mathcal{M}_0(U_0) = c]$$

where $U \sim \text{Unif}([d]_+)$ and $U_0 \sim \text{Unif}([d_0]_+)$. Then $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0} : \mathcal{X}^n \rightarrow \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$ has the following properties:

- i. $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ is $(\varepsilon_1 + \varepsilon_2 + \varepsilon_3, 0)$ -differentially private.
- ii. $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ has (a, β) -per-query accuracy.
- iii. $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ has $(a, \beta \cdot |\mathcal{X}|)$ -simultaneous accuracy.
- iv. $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ has running time

$$\tilde{O}(n \cdot \log d_0) + O(n) \cdot (\log d + \text{Time}(\mathcal{M}) + \text{Time}(\mathcal{S})) + O(n \log n \cdot \log |\mathcal{X}|)$$

where $\text{Time}(\mathcal{S})$ is the maximum time over $v \in [n]$ to sample from the distribution $\mathcal{S}_v \sim \text{Unif}(\{u_0 \in [d_0]_+ : \mathcal{M}_0(u_0) = v\})$.

- v. Given $h = \text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}(D)$, for all $x \in \mathcal{X}$ the count h_x can be evaluated in time

$$O(n) \cdot \tilde{O}(\log d_0) + \text{Time}(\mathcal{M}_0)$$

As discussed earlier, a natural choice for \mathcal{M}_0 is to take $\mathcal{M}_0(u) = \mathcal{M}(0, u)$ for all $u \in [d]$. However, d may not satisfy the required constraints. In the next section (see Lemma 8.4), we will show how to construct \mathcal{M}_0 for the counting query algorithms of Section 4 at only a constant factor loss in privacy (i.e. $\varepsilon_2 = O(\varepsilon)$ and $\varepsilon_3 = O(\varepsilon)$).

Proof of i. Let $D, D' \in \mathcal{X}^n$ be neighboring datasets. Let $A = \{x \in \mathcal{X} : c_x(D) > 0\}$. Similarly, define $A' = \{x \in \mathcal{X} : c_x(D') > 0\}$. Let $B = A \cup A'$. Notice that $|B| \leq n + 1$. Let $h \sim \text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}(D)$ and $h' \sim \text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}(D')$. Let $g \sim \text{Unif}(\mathcal{G}_{\mathcal{M}_0}(\mathcal{X}))$ and $r \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$. Then by construction

$$\begin{aligned} \Pr[h = r] &= \Pr[\forall x \in B \ h_x = r_x] \cdot \frac{\Pr[h = r \mid \forall x \in A \ h_x = r_x]}{\Pr[\forall x \in B \ h_x = r_x \mid \forall x \in A \ h_x = r_x]} \\ &= \frac{\Pr[\forall x \in B \ h_x = r_x]}{|\{g' \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X}) : \forall x \in A \ g'_x = r_x\}|} \cdot \frac{|\{g' \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X}) : \forall x \in A \ g'_x = r_x\}|}{|\{g' \in \mathcal{G}_{\mathcal{M}_0}(\mathcal{X}) : \forall x \in B \ g'_x = r_x\}|} \\ &= \Pr[\forall x \in B \ h_x = r_x] \cdot \Pr[g = r \mid \forall x \in B \ g_x = r_x] \end{aligned}$$

Now, because $|B \setminus A| \leq 1$ and $|B \setminus A'| \leq 1$ along with Lemma 8.1 Parts *i-ii*

$$\begin{aligned}
& \frac{\Pr[h = r]}{\Pr[g = r \mid \forall x \in B \ g_x = r_x]} \\
&= \left(\prod_{x \in A} \Pr[\mathcal{M}(c_x(D), U) = r_x] \right) \left(\prod_{x \in B \setminus A} \Pr[h_x = r_x \mid \forall x \in A \ h_x = r_x] \right) \\
&= \left(\prod_{x \in A} \Pr[\mathcal{M}(c_x(D), U) = r_x] \right) \left(\prod_{x \in B \setminus A} \Pr[g_x = r_x \mid \forall x \in A \ g_x = r_x] \right) \\
&= \left(\prod_{x \in A} \Pr[\mathcal{M}(c_x(D), U) = r_x] \right) \left(\prod_{x \in B \setminus A} \Pr[\mathcal{M}_0(U_0) = r_x] \right) \\
&\leq e^{\varepsilon_2} \cdot \prod_{x \in B} \Pr[\mathcal{M}(c_x(D), U) = r_x] \\
&\leq e^{\varepsilon_1 + \varepsilon_2} \cdot \prod_{x \in B} \Pr[\mathcal{M}(c_x(D'), U) = r_x] \\
&\leq \left(e^{\varepsilon_1 + \varepsilon_2} \cdot \prod_{x \in A'} \Pr[\mathcal{M}(c_x(D'), U) = r_x] \right) \left(e^{\varepsilon_3} \cdot \prod_{x \in B \setminus A'} \Pr[\mathcal{M}_0(U_0) = r_x] \right) \\
&= e^{\varepsilon_1 + \varepsilon_2 + \varepsilon_3} \cdot \frac{\Pr[h' = r]}{\Pr[g = r \mid \forall x \in B \ g_x = r_x]}
\end{aligned}$$

Therefore, $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ is $(\varepsilon_1 + \varepsilon_2 + \varepsilon_3, 0)$ -differentially private. \square

Proof of ii. Let $h \sim \text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}(D)$. Let $A = \{x \in \mathcal{X} : c_x(D) > 0\}$. If $x \in A$, then $h_x \sim \mathcal{M}(c_x(D), U)$ with accuracy following from \mathcal{M} having (a, β) -accuracy.

Otherwise, for $x \in \mathcal{X} \setminus A$, let $g \sim \text{Unif}(\mathcal{G}_{\mathcal{M}_0}(\mathcal{X}))$. Notice that $c_x(D) = 0$ and $|A| \leq n$. By construction and Lemma 8.1 Parts *i-ii*

$$\begin{aligned}
\Pr[|h_x| \leq a] &= \sum_{c \in [n]^A} \Pr[\forall x' \in A \ h_{x'} = c_{x'}] \cdot \Pr[h_x \leq a \mid \forall x' \in A \ h_{x'} = c_{x'}] \\
&= \sum_{c \in [n]^A} \Pr[\forall x' \in A \ h_{x'} = c_{x'}] \cdot \Pr[g_x \leq a \mid \forall x' \in A \ g_{x'} = c_{x'}] \\
&= \sum_{c \in [n]^A} \Pr[\forall x' \in A \ h_{x'} = c_{x'}] \cdot \Pr[g_x \leq a] \\
&= \Pr[g_x \leq a] \\
&= \Pr[\mathcal{M}_0(U_0) \leq a] \\
&\geq 1 - \beta
\end{aligned}$$

with the last inequality by assumption. Therefore, $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ has (a, β) -per-query accuracy. \square

Proof of iii. $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ has $(a, \beta \cdot |\mathcal{X}|)$ -simultaneous accuracy by a union bound over each $x \in \mathcal{X}$ along with the previous part. \square

Proof of iv-v. By Theorem 5.2 Part *iv* and Lemma 8.1 Parts *iii-iv*, we get the desired bounds on running time. \square

8.3 Constructing the Empty-Bin Sampler

The following lemma allows us to construct the empty-bin sampler satisfying the constraints of Theorem 8.3 for the counting query algorithms of Section 4.

Lemma 8.4. *Let deterministic $\mathcal{M} : [n] \times [d]_+ \rightarrow [n]$ such that $\mathcal{M}(0, u)$ is non-decreasing function in u . Let $d_0 \in \mathbb{N}_+$ such that $d_0 \geq (4/3) \cdot d$. Let $U \sim \text{Unif}([d]_+)$ and $U_0 \sim \text{Unif}([d_0]_+)$. Then there is a deterministic algorithm $\mathcal{M}_0 : [d_0]_+ \rightarrow [n]$ with the following properties:*

i. \mathcal{M}_0 is non-decreasing.

ii. For all $c \in [n]$

$$e^{-2 \cdot d/d_0} \cdot \Pr[\mathcal{M}(0, U) = c] \leq \Pr[\mathcal{M}_0(U_0) = c] \leq e^{d/d_0} \cdot \Pr[\mathcal{M}(0, U) = c]$$

Moreover, if d_0 is a multiple of d , then $\mathcal{M}_0(U_0) \sim \mathcal{M}(0, U)$.

iii. For all $a \in [n]$

$$\Pr[\mathcal{M}_0(U_0) \leq a] \geq \Pr[\mathcal{M}(0, U) \leq a]$$

In particular, if \mathcal{M} has (α, β) -accuracy, then $\Pr[\mathcal{M}_0(U_0) \leq a] \geq 1 - \beta$.

iv. \mathcal{M}_0 has running time

$$\tilde{O}(\log d_0) + \text{Time}(\mathcal{M})$$

v. For any $v \in [n]$, define the distribution $\mathcal{S}_v \sim \text{Unif}(\{u_0 \in [d_0]_+ : \mathcal{M}_0(u_0) = v\})$. Then \mathcal{S}_v can be sampled in time

$$\tilde{O}(\log d_0) + O(\text{Time}(F_{\mathcal{M}}))$$

Proof. (Construction) We start by considering a sample uniformly at random from $[d_0]_+$ and then map it to $[d]_+$ such that the resulting number is almost uniform while preserving monotonicity. The algorithm is as follows:

Algorithm 8.5. $\mathcal{M}_0(u_0)$ for $u_0 \in [d_0]_+$

1. Pick $q \in \mathbb{N}$ and $r \in [d - 1]$ such that $d_0 = q \cdot d + r$.
2. Define the function $f : [d_0]_+ \rightarrow [d]_+$

$$f(u_0) = \begin{cases} \lceil u_0 / (q + 1) \rceil & \text{if } r \neq 0 \text{ and } u_0 \leq r \cdot (q + 1) \\ \lceil (u_0 - r) / q \rceil & \text{if } r = 0 \text{ or } u_0 > r \cdot (q + 1) \end{cases}$$

3. Return $\mathcal{M}(0, f(u_0))$.

Proof of i. Notice that f as defined in Algorithm 8.5 is non-decreasing. Therefore, $\mathcal{M}_0(u_0) = \mathcal{M}(0, f(u_0))$ is non-decreasing in u_0 as $\mathcal{M}(0, u)$ is non-decreasing in u . \square

Proof of ii. Notice that $\lceil (u_0 - r)/q \rceil = \lceil (u_0 - r \cdot (q + 1))/q \rceil + r$. So $|\{u_0 \in [d_0]_+ : f(u_0) = u\}| = q$ or $q + 1$ for all $u \in [d]_+$. Thus,

$$\begin{aligned} \Pr[\mathcal{M}_0(U_0) = c] &\leq \frac{q+1}{d_0} \cdot |\{u \in [d]_+ : \mathcal{M}(0, u) = c\}| \\ &= \frac{(q+1) \cdot d}{d_0} \cdot \Pr[\mathcal{M}(0, U) = c] \\ &\leq \left(1 + \frac{d}{d_0}\right) \cdot \Pr[\mathcal{M}(0, U) = c] \\ &\leq e^{d/d_0} \cdot \Pr[\mathcal{M}(0, U) = c] \end{aligned}$$

and because $d/d_0 \leq 3/4$

$$\begin{aligned} \Pr[\mathcal{M}_0(U_0) = c] &\geq \frac{q \cdot d}{d_0} \cdot \Pr[\mathcal{M}(0, U) = c] \\ &\geq \left(1 - \frac{d}{d_0}\right) \cdot \Pr[\mathcal{M}(0, U) = c] \\ &\geq e^{-2 \cdot d/d_0} \cdot \Pr[\mathcal{M}(0, U) = c] \end{aligned}$$

If d_0 is a multiple of d , then $r = 0$ and $|\{u_0 \in [d_0]_+ : f(u_0) = u\}| = q$ for all $u \in [d]_+$. So $\Pr[\mathcal{M}_0(U_0) = c] = \Pr[\mathcal{M}(0, U) = c]$ for all $c \in [n]$. \square

Proof of iii. Let $u \in [d]_+$. Notice that

$$|\{u_0 \in [d_0]_+ : f(u_0) = u\}| = \begin{cases} q+1 & \text{if } u \leq r \\ q & \text{otherwise} \end{cases}$$

Then for all $a \in [n]$

$$\begin{aligned} \Pr[\mathcal{M}_0(U_0) \leq a] &= \frac{q \cdot d}{d_0} \cdot \Pr[\mathcal{M}(0, U) \leq a] + \frac{r}{d_0} \cdot \Pr[\mathcal{M}(0, U) \leq a \mid U \leq r] \\ &\geq \frac{q \cdot d}{d_0} \cdot \Pr[\mathcal{M}(0, U) \leq a] + \frac{r}{d_0} \cdot \Pr[\mathcal{M}(0, U) \leq a] \end{aligned}$$

as $\mathcal{M}(0, u)$ is non-decreasing in u . Therefore, $\Pr[\mathcal{M}_0(U_0) \leq a] \geq \Pr[\mathcal{M}(0, U) \leq a]$ for all $a \in [n]$. \square

Proof of iv. Evaluating f takes time $\tilde{O}(\log d_0)$. So we have the desired running time overall. \square

Proof of v. Because $\mathcal{M}(0, u)$ is non-decreasing in u and $f(u_0)$ is non-decreasing in u_0 , we have

$$\begin{aligned} \text{supp}(\mathcal{S}_v) &= \{u_0 \in [d_0]_+ : \mathcal{M}(0, f(u_0)) = v\} \\ &= \{u_0 \in [d_0]_+ : F_{\mathcal{M}}(v-1) < f(u_0) \leq F_{\mathcal{M}}(v)\} \\ &= \{\min\{u_0 \in [d_0]_+ : f(u_0) \geq F_{\mathcal{M}}(v-1) + 1\}, \dots, \max\{u_0 \in [d_0]_+ : f(u_0) \leq F_{\mathcal{M}}(v)\}\} \end{aligned}$$

One can show for $u \in [d]_+$ that

$$\max\{u_0 \in [d_0]_+ : f(u_0) \leq u\} = \begin{cases} (q+1) \cdot u & \text{if } r \neq 0 \text{ and } u \leq r \\ q \cdot u + r & \text{if } r = 0 \text{ or } u > r \end{cases}$$

and $\min\{u_0 \in [d_0]_+ : f(u_0) \geq u\} = 1 + \max\{u_0 \in [d_0]_+ : f(u_0) \leq u-1\}$. So both endpoints can be computed in time $\tilde{O}(\log d_0)$. \square

Now, we are ready to obtain a private histogram algorithm that achieves pure differential privacy, has the same accuracy guarantees (up to constant factors) as the Laplace mechanism and has running polynomial in n and $\log |\mathcal{X}|$.

Theorem 8.6. *Let $\varepsilon, \beta_0 \in \mathbb{N}^{-1}$ and $\mathcal{M} = \text{FastSample}_{n, \varepsilon', \gamma}$ where $\varepsilon' = 1/\lceil 10/(9\varepsilon) \rceil$ and $\gamma = \beta_0/(2|\mathcal{X}|)$. Then there exists deterministic $\mathcal{M}_0 : [d_0]_+ \rightarrow [n]$ with $\log d_0 = \tilde{O}(1/\varepsilon) \cdot (\log n + \log |\mathcal{X}| + \log(1/\beta_0))$ such that $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0} : \mathcal{X}^n \rightarrow \mathcal{G}_{\mathcal{M}_0}(\mathcal{X})$ has the following properties:*

i. $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ is $(\varepsilon, 0)$ -differentially private.

ii. For every $\beta \geq 2\gamma$, $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ has (a, β) -per-query accuracy for

$$a = \left\lceil \frac{5}{\varepsilon} \ln \left(\frac{2}{\beta} \right) \right\rceil$$

iii. For every $\beta \geq \beta_0$, $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ has (a, β) -per-query accuracy for

$$a = \left\lceil \frac{5}{\varepsilon} \ln \left(\frac{2 \cdot |\mathcal{X}|}{\beta} \right) \right\rceil$$

iv. $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ has running time

$$\tilde{O} \left(\frac{n}{\varepsilon} \cdot \log \frac{|\mathcal{X}|}{\beta_0} \right)$$

v. Given $h = \text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}(D)$, for all $x \in \mathcal{X}$ the count h_x can be evaluated in time

$$\tilde{O} \left(\frac{n}{\varepsilon} \cdot \log \frac{|\mathcal{X}|}{\beta_0} \right)$$

Proof. Let $d_0 = 2^{2 \cdot 3^\ell}$ where $\ell = \lceil \log_3(\lceil \log_2 \max\{|\mathcal{X}|, 30 \cdot d/\varepsilon\} \rceil / 2) \rceil$ where $\log d = \tilde{O}(1/\varepsilon) \cdot (\log n + \log |\mathcal{X}| + \log(1/\beta_0))$ as defined in Algorithm 4.11. Notice that $d_0 \geq |\mathcal{X}|$ and $d/d_0 \leq \varepsilon/30$. In addition, $d_0 = O(|\mathcal{X}|^3 + (d/\varepsilon)^3)$. So $\log d_0 = \tilde{O}(1/\varepsilon) \cdot (\log n + \log |\mathcal{X}| + \log(1/\beta_0))$. Now, let $\mathcal{M}_0 : [d_0]_+ \rightarrow [n]$ be defined as in Lemma 8.4 for \mathcal{M} . The proof follows from Theorem 4.10 and Theorem 8.3. \square

In the following figure, we show using CompactHistogram with FastSample is asymptotically faster than when using it with GeoSample , particularly in the case when $\varepsilon \gg 1/n$, with only a small constant loss in accuracy.

\mathcal{M}	Running Time	Evaluation Time
GeoSample	$\tilde{O}(n^2 \cdot \log(1/\varepsilon) + n \cdot \log \mathcal{X})$	$\tilde{O}(n^2 \cdot \log(1/\varepsilon) + n \cdot \log \mathcal{X})$
FastSample	$\tilde{O}((n/\varepsilon) \cdot \log(\mathcal{X} /\beta))$	$\tilde{O}((n/\varepsilon) \cdot \log(\mathcal{X} /\beta))$

\mathcal{M}	(a, β) -Per-Query	(a, β) -Simultaneous
GeoSample	$\left\lceil \frac{5}{\varepsilon} \ln \frac{1}{\beta} \right\rceil$	$\left\lceil \frac{5}{\varepsilon} \ln \frac{ \mathcal{X} }{\beta} \right\rceil$
FastSample	$\left\lceil \frac{5}{\varepsilon} \ln \frac{2}{\beta} \right\rceil$	$\left\lceil \frac{5}{\varepsilon} \ln \frac{2 \mathcal{X} }{\beta} \right\rceil$

Figure 6: The running time and errors of $\text{CompactHistogram}_{\mathcal{M}, \mathcal{M}_0}$ for the counting query algorithms of Section 4 using the empty-bin sampler \mathcal{M}_0 defined in Lemma 8.4 with $\log d_0 = O(\log(1/\varepsilon) + \log d + \log |\mathcal{X}|)$. For more details, see Theorem 8.6.

Acknowledgments

We thank the Harvard Privacy Tools differential privacy research group, particularly Mark Bun and Kobbi Nissim, for informative discussions and feedback. And we thank Ashwin Machanavajjhala, Frank McSherry, Uri Stemmer, and the anonymous TDPD and ITCS reviewers for their helpful comments.

References

- [BBKN14] Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. *Machine learning*, 94(3):401–437, 2014.
- [BLR13] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *J. ACM*, 60(2):12:1–12:25, may 2013. doi:10.1145/2450142.2450148.
- [BNS16] Mark Bun, Kobbi Nissim, and Uri Stemmer. Simultaneous private learning of multiple concepts. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, ITCS '16*, pages 369–380, New York, NY, USA, 2016. ACM. doi:10.1145/2840728.2840747.
- [BV18] Victor Balcer and Salil Vadhan. Differential privacy on finite computers. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 94. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [CDK17] Bryan Cai, Constantinos Daskalakis, and Gautam Kamath. Priv’it: Private and sample efficient identity testing. *CoRR*, abs/1703.10127, 2017. arXiv:1703.10127.
- [CKKL12] Mahdi Cheraghchi, Adam Klivans, Pravesh Kothari, and Homin K. Lee. Submodular functions are noise stable. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12*, pages 1586–1592, Philadelphia, PA, USA, 2012. Society for Industrial and Applied Mathematics.

- [CLRS09] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2009.
- [CPST12] Graham Cormode, Cecilia Procopiuc, Divesh Srivastava, and Thanh T. L. Tran. Differentially private summaries for sparse data. In *Proceedings of the 15th International Conference on Database Theory, ICDT '12*, pages 299–311, New York, NY, USA, 2012. ACM. doi:10.1145/2274576.2274608.
- [CTUW14] Karthekeyan Chandrasekaran, Justin Thaler, Jonathan Ullman, and Andrew Wan. Faster private release of marginals on small databases. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science, ITCS '14*, pages 387–402, New York, NY, USA, 2014. ACM. doi:10.1145/2554797.2554833.
- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Eurocrypt*, volume 4004, pages 486–503. Springer, 2006.
- [DL09] Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing, STOC '09*, pages 371–380, New York, NY, USA, 2009. ACM. doi:10.1145/1536414.1536466.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, volume 3876, pages 265–284. Springer, 2006.
- [DNT15] Cynthia Dwork, Aleksandar Nikolov, and Kunal Talwar. Efficient algorithms for privately releasing marginals via convex relaxations. volume 53, pages 650–673. Springer, 2015.
- [GMP13] Ivan Gazeau, Dale Miller, and Catuscia Palamidessi. Preserving differential privacy under finite-precision semantics. In *Proceedings 11th International Workshop on Quantitative Aspects of Programming Languages and Systems, QAPL 2013, Rome, Italy, March 23-24, 2013.*, pages 1–18, 2013. doi:10.4204/EPTCS.117.1.
- [GRS12] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. *SIAM Journal on Computing*, 41(6):1673–1693, 2012.
- [GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. *Theory of Cryptography*, pages 339–356, 2012.
- [HRS12] Moritz Hardt, Guy N. Rothblum, and Rocco A. Servedio. Private data release via learning thresholds. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12*, pages 168–187, Philadelphia, PA, USA, 2012. Society for Industrial and Applied Mathematics.
- [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. doi:10.1145/1806689.1806786.
- [KKMN09] Aleksandra Korolova, Krishnaram Kenthapadi, Nina Mishra, and Alexandros Ntoulas. Releasing search queries and clicks privately. In *Proceedings of the 18th International Conference on World Wide Web, WWW '09*, pages 171–180, New York, NY, USA, 2009. ACM. doi:10.1145/1526709.1526733.

- [KLN⁺11] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- [Lin99] Jacobus Hendricus van Lint. Introduction to coding theory. 1999.
- [Mir12] Ilya Mironov. On significance of the least significant bits for differential privacy. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security, CCS '12*, pages 650–661, New York, NY, USA, 2012. ACM. doi:10.1145/2382196.2382264.
- [TUV12] Justin Thaler, Jonathan Ullman, and Salil Vadhan. Faster algorithms for privately releasing marginals. In *International Colloquium on Automata, Languages, and Programming*, pages 810–821. Springer, 2012.
- [UV11] Jonathan Ullman and Salil P. Vadhan. PCPs and the hardness of generating private synthetic data. In *TCC*, volume 6597, pages 400–416. Springer, 2011.
- [vzGG13] Joachim von zur Gathen and Jürgen Gerhard. *Modern computer algebra*. Cambridge university press, 2013.

A Efficient Sampling of Distinct Elements

In this section we show how to efficiently sample distinct elements from a subset of \mathcal{X} .

Lemma A.1. *There exists an algorithm that given an integer m specifying the set $\mathcal{X} = [m]_+$ and a subset $A \subseteq \mathcal{X}$, samples a uniformly random sequence of $r \leq |\mathcal{X} \setminus A|$ distinct elements from $\mathcal{X} \setminus A$ with running time*

$$O(|A| \log |A| \cdot \log |\mathcal{X}|) + O(r \cdot \log^2 |\mathcal{X}| \cdot \log(|A| + r))$$

To prove this lemma, we will use a data structure that supports efficiently computing the numbers of elements in the tree less than a given value. We will use this data structure to store the elements in \mathcal{X} which we do not want to sample.

Proposition A.2 (Order-Statistic Tree [CLRS09] Chapter 14.1). *There exists a data structure T maintaining a set over \mathcal{X} (let $|T|$ denote the size of the set maintained by T) with the following properties:*

- i. Inserting an element $x \in \mathcal{X}$ into T takes $O(\log |T| \cdot \log |\mathcal{X}|)$ time.*
- ii. T can be represented in $O(|T| \cdot \log |\mathcal{X}|)$ bits.*
- iii. For all $x \in \mathcal{X}$, the quantity $|\{x' \in T : x' \leq x\}|$ can be computed in $O(\log |T| \cdot \log |\mathcal{X}|)$ time.*

Proof of Lemma A.1. We define the sampling algorithm as follows:

Algorithm A.3. $\text{DistinctSample}(A, r)$ for $A \subseteq \mathcal{X}$ and $r \in \mathbb{N}_+$ such that $r \leq |\mathcal{X} \setminus A|$

1. Let $S = \emptyset$ and T be the data structure defined in Proposition A.2.
 2. For $x \in A$, insert x into T .
 3. For $i \in [r]_+$:
 - (a) Let $m' = |\mathcal{X}| - |T|$.
 - (b) Sample z uniformly at random from $[m']_+$.
 - (c) Perform binary search over \mathcal{X} to find $s = \min\{x \in \mathcal{X} : x - |\{x' \in T : x' \leq x\}| = z\}$.
 - (d) Insert s into T and let $S = S \cup \{s\}$.
 4. Return S
-

We prove correctness by induction on r . We start with the base case $r = 1$. Notice that $x - |\{x' \in T : x' \leq x\}| = |\{x' \in \mathcal{X} \setminus A : x' \leq x\}|$. Thus for all $z \in [m']_+$ there exists $x \in \mathcal{X} \setminus A$ such that $|\{x' \in \mathcal{X} \setminus A : x' \leq x\}| = z$.

Now, we show $s \in \mathcal{X} \setminus A$. Let $x \in A$. If $|\{x' \in \mathcal{X} \setminus A : x' \leq x\}| = 0$, then $s \neq x$ as $z \geq 1$. Otherwise $|\{x' \in \mathcal{X} \setminus A : x' \leq x\}| = |\{x' \in \mathcal{X} \setminus A : x' \leq (x - 1)\}|$ which also implies $s \neq x$ by definition of s . Therefore, $s \in \mathcal{X} \setminus A$.

Notice that two different values of z cannot output the same s and $m' = |\mathcal{X} \setminus A|$. Therefore, $\text{DistinctSample}(A, 1)$ is uniformly distributed over $\mathcal{X} \setminus A$.

For the induction step, let $S \sim \text{DistinctSample}(A, 1)$ and assume $\text{DistinctSample}(A', r - 1)$ is uniformly distributed over random sequences of $r - 1$ elements from $\mathcal{X} \setminus A'$ for any $A' \subseteq \mathcal{X}$. Then

$$\text{DistinctSample}(A, r) \sim S \cup \text{DistinctSample}(A \cup S, r - 1)$$

Therefore $\text{DistinctSample}(A, r)$ is uniformly distributed over random sequences of r elements from $\mathcal{X} \setminus A$.

Now, we analyze the running time of DistinctSample . Step 2 can be done in $O(|A| \log |A| \cdot \log |\mathcal{X}|)$ time by Proposition A.2 Part *i*. Each of the r iterations of step 3 is dominated by step 3c which takes $O(\log^2 |\mathcal{X}| \cdot \log(|A| + r))$ time by Proposition A.2 Part *iii* as $|T| \leq |A| + r$. \square