# Boosting the Accuracy of Differentially Private Histograms Through Consistency 

Michael Hay<br>Vibhor Rastogi<br>Gerome Miklau, University of Massachusetts - Amherst<br>Dan Suciu

# Boosting the Accuracy of Differentially Private Histograms Through Consistency 

Michael $\mathrm{Hay}^{\dagger}$, Vibhor Rastogi ${ }^{\ddagger}$, Gerome Miklau ${ }^{\dagger}$, Dan Suciu ${ }^{\ddagger}$<br>$\dagger$ University of Massachusetts Amherst<br>\{mhay,miklau\}@cs.umass.edu<br>$\ddagger$ University of Washington<br>\{vibhor,suciu\}@cs.washington.edu


#### Abstract

We show that it is possible to significantly improve the accuracy of a general class of histogram queries while satisfying differential privacy. Our approach carefully chooses a set of queries to evaluate, and then exploits consistency constraints that should hold over the noisy output. In a postprocessing phase, we compute the consistent input most likely to have produced the noisy output. The final output is differentially-private and consistent, but in addition, it is often much more accurate. We show, both theoretically and experimentally, that these techniques can be used for estimating the degree sequence of a graph very precisely, and for computing a histogram that can support arbitrary range queries accurately.


## 1. INTRODUCTION

Recent work in differential privacy [9] has shown that it is possible to analyze sensitive data while ensuring strong privacy guarantees. Differential privacy is typically achieved through random perturbation: the analyst issues a query and receives a noisy answer. To ensure privacy, the noise is carefully calibrated to the sensitivity of the query. Informally, query sensitivity measures how much a small change to the database - such as adding or removing a person's private record-can affect the query answer. Such query mechanisms are simple, efficient, and often quite accurate. In fact, one mechanism has recently been shown to be optimal for a single counting query [10]-i.e., there is no better noisy answer to return under the desired privacy objective.

However, analysts typically need to compute multiple statistics on a database. Differentially private algorithms extend nicely to a set of queries, but there can be difficult trade-offs among alternative strategies for answering a workload of queries. Consider the analyst of a private student database who requires answers to the following queries: the total number of students, $x_{t}$, the number of students $x_{A}$, $x_{B}, x_{C}, x_{D}, x_{F}$ receiving grades $\mathrm{A}, \mathrm{B}, \mathrm{C}, \mathrm{D}$, and F respectively, and the number of passing students, $x_{p}$ (grade D or


Figure 1: Our approach to querying private data.
higher).
Using a differentially private interface, a first alternative is to request noisy answers for just $\left(x_{A}, x_{B}, x_{C}, x_{D}, x_{F}\right)$ and use those answers to compute answers for $x_{t}$ and $x_{p}$ by summation. The sensitivity of this set of queries is 1 because adding or removing one tuple changes exactly one of the five outputs by a value of one. Therefore, the noise added to individual answers is low and the noisy answers are accurate estimates of the truth. Unfortunately, the noise accumulates under summation, so the estimates for $x_{t}$ and $x_{p}$ are worse.

A second alternative is to request noisy answers for all queries ( $x_{t}, x_{p}, x_{A}, x_{B}, x_{C}, x_{D}, x_{F}$ ). This query set has sensitivity 3 (one change could affect three return values, each by a value of one), and the privacy mechanism must add more noise to each component. This means the estimates for $x_{A}, x_{B}, x_{C}, x_{D}, x_{F}$ are worse than above, but the estimates for $x_{t}$ and $x_{p}$ may be more accurate. There is another concern, however: inconsistency. The noisy answers are likely to violate the following constraints, which one would naturally expect to hold: $x_{t}=x_{p}+x_{F}$ and $x_{p}=x_{A}+x_{B}+x_{C}+x_{D}$. This means the analyst must find a way to reconcile the fact that there are two different estimates for the total number of students and two different estimates for the number of passing students. We propose a technique for resolving inconsistency in a set of noisy answers, and show that doing so can actually increase accuracy. As a result, we show that strategies inspired by the second alternative can be superior in many cases.

Overview of Approach. Our approach, shown pictorially in Figure 1, involves three steps.

First, given a task - such as computing a histogram over student grades-the analyst chooses a set of queries $\mathbf{Q}$ to send to the data owner. The choice of queries will depend on the particular task, but in this work they are chosen so that

(a) Trace data
\(\left.\begin{array}{lll}Query Definitions: \& \mathbf{L}:\left\langle C_{000}, C_{001}, C_{010}, C_{011}\right\rangle <br>
\& \mathbf{H}:\left\langle C_{0 * *}, C_{00 *}, C_{01 *}, C_{000}, C_{001}, C_{010}, C_{011}\right\rangle <br>
\& \mathbf{S}: \operatorname{sort}(\mathbf{L}) \& <br>

\& \& Private output\end{array}\right)\) Inferred answer | True answer | $\tilde{\mathbf{L}}(I)=\langle 3,1,11,1\rangle$ |
| :--- | :--- |
| $\mathbf{L}(I)=\langle 2,0,10,2\rangle$ | $\tilde{\mathbf{H}}(I)=\langle 13,3,11,4,1,12,1\rangle$ |
| $\mathbf{H}(I)=\langle 14,2,12,2,0,10,2\rangle$ | $\overline{\mathbf{H}}(I)=\langle 14,3,11,3,0,11,0\rangle$ |
| $\mathbf{S}(I)=\langle 0,2,2,10\rangle$ | $\tilde{\mathbf{S}}(I)=\langle 1,2,0,11\rangle$ |

(b) Query variations

Figure 2: (a) Illustration of sample data representing a bipartite graph of network connections; (b) Definitions and sample values for alternative query sequences: $L$ counts the number of connections for each source, $H$ provides a hierarchy of range counts, and $S$ returns an ordered degree sequence for the implied graph.
constraints hold among the answers. For example, rather than issue ( $x_{A}, x_{B}, x_{C}, x_{D}, x_{F}$ ), the analyst would formulate the query as $\left(x_{t}, x_{p}, x_{A}, x_{B}, x_{C}, x_{D}, x_{F}\right)$, which has consistency constraints. The query set $\mathbf{Q}$ is sent to the data owner.

In the second step, the data owner answers the set of queries, using a standard differentially-private mechanism [9], as follows. The queries are evaluated on the private database and the true answer $\mathbf{Q}(I)$ is computed. Then random independent noise is added to each answer in the set, where the data owner scales the noise based on the sensitivity of the query set. The set of noisy answers $\tilde{q}$ is sent to the analyst. Importantly, because this step is unchanged from [9], it offers the same differential privacy guarantee.

The above step ensures privacy, but the set of noisy answers returned may be inconsistent. In the third and final step, the analyst post-processes the set of noisy answers to resolve inconsistencies among them. We propose a novel approach for resolving inconsistencies, called constrained inference, that finds a new set of answers $\bar{q}$ that is the "closest" set to $\tilde{q}$ that also satisfies the consistency constraints.

For two histogram tasks, our main technical contributions are efficient techniques for the third step and a theoretical and empirical analysis of the accuracy of $\bar{q}$. The surprising finding is that $\bar{q}$ can be significantly more accurate than $\tilde{q}$.

We emphasize that the constrained inference step has no impact on the differential privacy guarantee. The analyst performs this step without access to the private data, using only the constraints and the noisy answers, $\tilde{q}$. The noisy answers $\tilde{q}$ are the output of a differentially private mechanism; any post-processing of the answers cannot diminish this rigorous privacy guarantee. The constraints are properties of the query, not the database, and therefore known by the analyst a priori. For example, the constraint $x_{p}=x_{A}+x_{B}+x_{C}+x_{D}$ is simply the definition of $x_{p}$.

Intuitively, however, it would seem that if noise is added for privacy and then constrained inference reduces the noise, some privacy has been lost. In fact, our results show that existing techniques add more noise than is strictly necessary to ensure differential privacy. The extra noise provides no quantifiable gain in privacy but does have a significant cost in accuracy. We show that constrained inference can be an effective strategy for boosting accuracy.

The increase in accuracy we achieve depends on the input database and the privacy parameters. For instance, for some databases and levels of noise the perturbation may tend to
produce answers that do not violate the constraints. In this case the inference step would not improve accuracy. But we show that our inference process never reduces accuracy and give conditions under which it will boost accuracy. In practice, we find that many real datasets have data distributions for which our techniques significantly improve accuracy.

Histogram tasks. We demonstrate this technique on two specific tasks related to histograms. For relational schema $R(A, B, \ldots)$, we choose one attribute $A$ on which histograms are built (called the range attribute). We assume the domain of $A$, dom, is ordered.

We explain these tasks using sample data that will serve as a running example throughout the paper, and is also the basis of later experiments. The relation $R(s r c, d s t)$, shown in Fig. 2, represents a trace of network communications between a source IP address ( $s r c$ ) and a destination IP address (dst). It is bipartite because it represents flows through a gateway router from internal to external addresses.

In a conventional histogram, we form disjoint intervals for the range attribute and compute counting queries for each specified range. In our example, we use $s r c$ as the range attribute. There are four source addresses present in the table. If we ask for counts of all unit-length ranges, then the histogram is simply the sequence $\langle 2,0,10,2\rangle$ corresponding to the (out) degrees of the source addresses $\langle 000,001,010,011\rangle$.

Our first histogram task is an unattributed histogram, in which the intervals themselves are irrelevant to the analysis and so we report only a multiset of frequencies. For the example histogram, the multiset is $\{0,2,2,10\}$. An important instance of an unattributed histogram is the degree sequence of a graph, a crucial measure that is widely studied [18]. If the tuples of $R$ represent queries submitted to a search engine, and $A$ is the search term, then an unattributed histogram shows the frequency of occurrence of all terms (but not the terms themselves), and can be used to study the distribution.

For our second histogram task, we consider more conventional sequences of counting queries in which the intervals studied may be irregular and overlapping. In this case, simply returning unattributed counts is insufficient. And because we cannot predict ahead of time all the ranges of interest, our goal is to compute privately a set of statistics sufficient to support arbitrary interval counts and thus any histogram. We call this a universal histogram.

Continuing the example, a universal histogram allows the analyst to count the number of packets sent from any single address (e.g., the counts from source address 010 is 10 ), or from any range of addresses (e.g., the total number of packets is 14 , and the number of packets from a source address matching prefix $01 *$ is 12 ).

While a universal histogram can be used compute an unattributed histogram, we distinguish between the two because we show the latter can be computed much more accurately.

Contributions. For both unattributed and universal histograms, we propose a strategy for boosting the accuracy of existing differentially private algorithms. For each task, (1) we show that there is an efficiently-computable, closedform expression for the consistent query answer closest to a private randomized output; (2) we prove bounds on the error of the inferred output, showing under what conditions inference boosts accuracy; (3) we demonstrate significant improvements in accuracy through experiments on real data sets. Unattributed histograms are extremely accurate, with error at least an order of magnitude lower than existing techniques. Our approach to universal histograms can reduce error for larger ranges by 45-98\%, and improves on all ranges in some cases.

## 2. BACKGROUND

In this section, we introduce the concept of query sequences and how they can be used to support histograms. Then we review differential privacy and show how queries can be answered under differential privacy. Finally, we formalize our constrained inference process.

All of the tasks considered in this paper are formulated as query sequences where each element of the sequence is a simple count query on a range. We write intervals as $[x, y]$ for $x, y \in d o m$, and abbreviate interval $[x, x]$ as $[x]$. A counting query on range attribute $A$ is:

$$
c([x, y])=\text { Select count }(*) \text { From } \mathrm{R} \text { Where } x \leq \mathrm{R} \cdot \mathrm{~A} \leq y
$$

We use $\mathbf{Q}$ to denote a generic query sequence (please see Appendix A for an overview of notational conventions). When $\mathbf{Q}$ is evaluated on a database instance $I$, the output, $\mathbf{Q}(I)$, includes one answer to each counting query, so $\mathbf{Q}(I)$ is a vector of non-negative integers. The $i^{t h}$ query in $\mathbf{Q}$ is $\mathbf{Q}[i]$.

We consider the common case of a histogram over unitlength ranges. The conventional strategy is to simply compute counts for all unit-length ranges. This query sequence is denoted $\mathbf{L}$ :

$$
\mathbf{L}=\left\langle c\left(\left[x_{1}\right]\right), \ldots c\left(\left[x_{n}\right]\right)\right\rangle, x_{i} \in \operatorname{dom}
$$

Example 1. Using the example in Fig 2, we assume the domain of src contains just the 4 addresses shown. Query $\mathbf{L}$ is $\langle c([000]), c([001]), c([010]), c([011])\rangle$ and $\mathbf{L}(I)=\langle 2,0,10,2\rangle$.

### 2.1 Differential Privacy

Informally, an algorithm is differentially private if it is insensitive to small changes in the input. Formally, for any input database $I$, let $n b r s(I)$ denote the set of neighboring databases, each differing from $I$ by at most one record; i.e., if $I^{\prime} \in \operatorname{nbrs}(I)$, then $\left|\left(I-I^{\prime}\right) \cup\left(I^{\prime}-I\right)\right|=1$.

Definition 2.1 ( $\epsilon$-Differential privacy). Algorithm $A$ is $\epsilon$-differentially private if for all instances $I$, any $I^{\prime} \in$
nbrs(I), and any subset of outputs $S \subseteq \operatorname{Range}(A)$, the following holds:

$$
\operatorname{Pr}[A(I) \in S] \leq \exp (\epsilon) \times \operatorname{Pr}\left[A\left(I^{\prime}\right) \in S\right]
$$

where the probability is taken over the randomness of the $A$.
Differential privacy has been defined inconsistently in the literature. The original concept, called $\epsilon$-indistinguishability [9], defines neighboring databases using hamming distance rather than symmetric difference (i.e., $I^{\prime}$ is obtained from $I$ by replacing a tuple rather than adding/removing a tuple). The choice of definition affects the calculation of query sensitivity. We use the above definition (from Dwork [7]) but observe that our results also hold under indistinguishability, due to the fact that $\epsilon$-differential privacy (as defined above) implies $2 \epsilon$-indistinguishability.

To answer queries under differential privacy, we use the Laplace mechanism [9], which achieves differential privacy by adding noise to query answers, where the noise is sampled from the Laplace distribution. The magnitude of the noise depends on the query's sensitivity, defined as follows (adapting the definition to the query sequences considered in this paper).

Definition 2.2 (Sensitivity). Let $\mathbf{Q}$ be a sequence of counting queries. The sensitivity of $\mathbf{Q}$, denoted $S_{\mathbf{Q}}$, is

$$
\Delta_{\mathbf{Q}}=\max _{I, I^{\prime} \in n b r s(I)}\left\|\mathbf{Q}(I)-\mathbf{Q}\left(I^{\prime}\right)\right\|_{1}
$$

Throughout the paper, we use $\|\mathbf{X}-\mathbf{Y}\|_{p}$ to denote the $L_{p}$ distance between vectors $\mathbf{X}$ and $\mathbf{Y}$.

Example 2. The sensitivity of query $\mathbf{L}$ is 1. Database $I^{\prime}$ can be obtained from I by adding or removing a single record, which affects exactly one of the counts in $\mathbf{L}$ by exactly 1.

Given query $\mathbf{Q}$, the Laplace mechanism first computes the query answer $\mathbf{Q}(I)$ and then adds random noise independently to each answer. The noise is drawn from a zero-mean Laplace distribution with scale $\sigma$. As the following proposition shows, differential privacy is achieved if the Laplace noise is scaled appropriately to the sensitivity of $\mathbf{Q}$ and the privacy parameter $\epsilon$.

Proposition 1 (Laplace mechanism [9]). Let $\mathbf{Q}$ be a query sequence of length $d$, and let $\langle\operatorname{Lap}(\sigma)\rangle^{d}$ denote a $d$-length vector of i.i.d. samples from a Laplace with scale $\sigma$. The randomized algorithm $\tilde{\mathbf{Q}}$ that takes as input database $I$ and outputs the following vector is $\epsilon$-differentially private:

$$
\tilde{\mathbf{Q}}(I)=\mathbf{Q}(I)+\left\langle\operatorname{Lap}\left(\Delta_{\mathbf{Q}} / \epsilon\right)\right\rangle^{d}
$$

We apply this technique to the query $\mathbf{L}$. Since, $\mathbf{L}$ has sensitivity 1 , the following algorithm is $\epsilon$-differentially private:

$$
\tilde{\mathbf{L}}(I)=\mathbf{L}(I)+\langle\operatorname{Lap}(1 / \epsilon)\rangle^{n}
$$

We rely on Proposition 1 to ensure privacy for the query sequences we propose in this paper. We emphasize that the proposition holds for any query sequence $\mathbf{Q}$, regardless of correlations or constraints among the queries in $\mathbf{Q}$. Such dependencies are accounted for in the calculation of sensitivity. (For example, consider the correlated sequence $\mathbf{Q}$ that consists of the same query repeated $k$ times, then the sensitivity of $\mathbf{Q}$ is $k$ times the sensitivity of the query.)

We present the case where the analyst issues a single query sequence $\mathbf{Q}$. To support multiple query sequences, the protocol that computes a $\epsilon_{i}$-differentially private response to the $i^{t h}$ sequence is ( $\sum \epsilon_{i}$ )-differentially private.

To analyze the accuracy of the randomized query sequences proposed in this paper we quantify their error. $\tilde{\mathbf{Q}}$ can be considered an estimator for the true value $\mathbf{Q}(I)$. We use the common Mean Squared Error as a measure of accuracy.

Definition 2.3 (Error). For a randomized query sequence $\tilde{\mathbf{Q}}$ whose input is $\mathbf{Q}(I)$, the error $(\tilde{\mathbf{Q}})$ is $\sum_{i} \mathbb{E}(\tilde{\mathbf{Q}}[i]-$ $\mathbf{Q}[i])^{2}$ Here $\mathbb{E}$ is the expectation taken over the possible randomness in generating $\tilde{\mathbf{Q}}$.
For example, $\operatorname{error}(\tilde{\mathbf{L}})=\sum_{i} \mathbb{E}(\tilde{\mathbf{L}}[i]-\mathbf{L}[i])^{2}$ which simplifies to: $n \mathbb{E}\left[\operatorname{Lap}(1 / \epsilon)^{2}\right]=n \operatorname{Var}(\operatorname{Lap}(1 / \epsilon))=2 n / \epsilon^{2}$.

### 2.2 Constrained Inference

While $\tilde{\mathbf{L}}$ can be used to support unattributed and universal histograms under differential privacy, the main contribution of this paper is the development of more accurate query strategies based on the idea of constrained inference. The specific strategies are described in the next sections. Here, we formulate the constrained inference problem.

Given a query sequence $\mathbf{Q}$, let $\gamma_{\mathbf{Q}}$ denote a set of constraints which must hold among the (true) answers. The constrained inference process takes the randomized output of the query, denoted $\tilde{q}=\tilde{\mathbf{Q}}(I)$, and finds the sequence of query answers $\bar{q}$ that is "closest" to $\tilde{q}$ and also satisfies the constraints of $\gamma_{\mathbf{Q}}$. Here closest is determined by $L_{2}$ distance, and the result is the minimum $L_{2}$ solution:

Definition 2.4 (Minimum $L_{2}$ solution). Let $\mathbf{Q}$ be a query sequence with constraints $\gamma_{\mathbf{Q}}$. Given a noisy query sequence $\tilde{q}=\tilde{\mathbf{Q}}(I)$, a minimum $L_{2}$ solution, denoted $\bar{q}$, is a vector $\bar{q}$ that satisfies the constraints $\gamma_{\mathbf{Q}}$ and at the same time minimizes $\|\tilde{q}-\bar{q}\|_{2}$.

We use $\overline{\mathbf{Q}}$ to denote the two step randomized process in which the data owner first computes $\tilde{q}=\tilde{\mathbf{Q}}(I)$ and then computes the minimum $L_{2}$ solution from $\tilde{q}$ and $\gamma_{\mathbf{Q}}$. (Alternatively, the data owner can release $\tilde{q}$ and the latter step can be done by the analyst.) Importantly, the inference step has no impact on privacy, as stated below. (Proofs appear in the Appendix.)

Proposition 2. If $\tilde{\mathbf{Q}}$ satisfies $\epsilon$-differential privacy, then $\overline{\mathbf{Q}}$ satisfies $\epsilon$-differential privacy.

## 3. UNATTRIBUTED HISTOGRAMS

To support unattributed histograms, one could use the query sequence $\mathbf{L}$. However, it contains "extra" informationthe attribution of each count to a particular range - which is irrelevant for an unattributed histogram. Since the association between $\mathbf{L}[i]$ and $i$ is not required, any permutation of the unit-length counts is a correct response for the unattributed histogram. We formulate an alternative query that asks for the counts of $\mathbf{L}$ in rank order. As we will show, ordering does not increase sensitivity, but it does introduce inequality constraints that can be exploited by inference.

Formally, let $a_{i}$ refer to the answer to $\mathbf{L}[i]$ and let $\mathcal{U}=$ $\left\{a_{1}, \ldots, a_{n}\right\}$ be the multiset of answers to queries in $\mathbf{Q}$. We write $\operatorname{ran}_{i}(\mathcal{U})$ to refer to the $i^{\text {th }}$ smallest answer in $\mathcal{U}$. Then the query $\mathbf{S}$ is defined as

$$
\mathbf{S}=\left\langle\operatorname{rank}_{1}(\mathcal{U}), \ldots, \operatorname{rank}_{n}(\mathcal{U})\right\rangle
$$

Example 3. In the example in Fig 2, we have $\mathbf{L}(I)=$ $\langle 2,0,10,2\rangle$ while $\mathbf{S}(I)=\langle 0,2,2,10\rangle$. Thus, the answer $\mathbf{S}(I)$ contains the same counts as $\mathbf{L}(I)$ but in sorted order.

To answer $\mathbf{S}$ under differential privacy, we must determine its sensitivity.

Proposition 3. The sensitivity of $\mathbf{S}$ is 1 .
By Propositions 1 and 3, the following algorithm is $\epsilon$ differentially private:

$$
\tilde{\mathbf{S}}(I)=\mathbf{S}(I)+\langle\operatorname{Lap}(1 / \epsilon)\rangle^{n}
$$

Since the same magnitude of noise is added to $\mathbf{S}$ as to $\mathbf{L}$, the accuracy of $\tilde{\mathbf{S}}$ and $\tilde{\mathbf{L}}$ is the same. However, $\mathbf{S}$ implies a powerful set of constraints. Notice that the ordering occurs before noise is added. Thus, the analyst knows that the returned counts are ordered according to the true rank order. If the returned answer contains out-of-order counts, this must be caused by the addition of random noise, and they are inconsistent. Let $\gamma_{\mathbf{S}}$ denote the set of inequalities $\mathbf{S}[i] \leq \mathbf{S}[i+1]$ for $1 \leq i<n$. We show next how to exploit these constraints to boost accuracy.

### 3.1 Constrained Inference: Computing $\overline{\mathbf{S}}$

As outlined in the introduction, the analyst sends query $\mathbf{S}$ to the data owner and receives a noisy answer $\tilde{s}=\tilde{\mathbf{S}}(I)$, the output of the differentially private algorithm $\tilde{\mathbf{S}}$ evaluated on the private database $I$. We now describe a technique for post-processing $\tilde{s}$ to find an answer that is consistent with the ordering constraints.

Formally, given $\tilde{s}$, the objective is to find an $\bar{s}$ that minimizes $\|\tilde{s}-\bar{s}\|_{2}$ subject to the constraints $\bar{s}[i] \leq \bar{s}[i+1]$ for $1 \leq i<n$. The solution has a surprisingly elegant closedform. Let $\tilde{s}[i, j]$ be the subsequence of $j-i+1$ elements: $\langle\tilde{s}[i], \tilde{s}[i+1], \ldots, \tilde{s}[j]\rangle$. Let $\tilde{M}[i, j]$ be the mean of these elements, i.e. $\tilde{M}[i, j]=\sum_{k=i}^{j} \tilde{s}[k] /(j-i+1)$.

Theorem 1. Denote $L_{k}=\min _{j \in[k, n]} \max _{i \in[1, j]} \tilde{M}[i, j]$ and $U_{k}=\max _{i \in[1, k]} \min _{j \in[i, n]} \tilde{M}[i, j]$. The minimum $L_{2}$ solution $\bar{s}$, is unique and given by: $\bar{s}[k]=L_{k}=U_{k}$.

Since we first stated this result in a technical report [13], we have learned that this problem is an instance of isotonic regression (i.e., least squares regression under ordering constraints on the estimands). The statistics literature gives several characterizations of the solution, including the above min-max formulas (cf. Barlow et al. [3]), as well as linear time algorithms for computing it (cf. Barlow et al. [2]).

Example 4. We give three examples of $\tilde{s}$ and its closest ordered sequence $\bar{s}$. First, suppose $\tilde{s}=\langle 9,10,14\rangle$. Since $\tilde{s}$ is already ordered, $\bar{s}=\tilde{s}$. Second, $\tilde{s}=\langle 9,14,10\rangle$, where the last two elements are out of order. The closest ordered sequence is $\bar{s}=\langle 9,12,12\rangle$. Finally, let $\tilde{s}=\langle 14,9,10,15\rangle$. The sequence is in order except for $\tilde{s}[1]$. While changing the first element from 14 to 9 would make it ordered, its distance from $\tilde{s}$ would be $(14-9)^{2}=25$. In contrast, $\bar{s}=$ $\langle 11,11,11,15\rangle$ and $\|\tilde{s}-\bar{s}\|_{2}=14$.

### 3.2 Utility Analysis: the Accuracy of $\bar{S}$

Prior work in isotonic regression has shown inference cannot hurt, i.e., the accuracy of $\overline{\mathbf{S}}$ is no lower than $\tilde{\mathbf{S}}$ [14].


Figure 3: Example of how $\bar{s}$ reduces the error of $\tilde{s}$.
However, we are not aware of any results that give conditions for which $\overline{\mathbf{S}}$ is more accurate than $\tilde{\mathbf{S}}$. Before presenting a theoretical statement of such conditions, we first give an illustrative example.

Example 5. Figure 3 shows a sequence $\mathbf{S}(I)$ along with a sampled $\tilde{s}$ and inferred $\bar{s}$. While the values in $\tilde{s}$ deviate considerably from $\mathbf{S}(I)$, $\bar{s}$ lies very close to the true answer. In particular, for subsequence $[1,20]$, the true sequence $\mathbf{S}(I)$ is uniform and the constrained inference process effectively averages out the noise of $\tilde{s}$. At the twenty-first position, which is a unique count in $\mathbf{S}(I)$, and constrained inference does not refine the noisy answer, i.e., $\bar{s}[21]=\tilde{s}[21]$.

Fig 3 suggests that error $(\overline{\mathbf{S}})$ will be low for sequences in which many counts are the same (Fig 7 in Appendix C gives another intuitive view of the error reduction). The following theorem quantifies the accuracy of $\overline{\mathbf{S}}$ precisely. Let $n$ and $d$ denote the number of values and the number of distinct values in $\mathbf{S}(I)$ respectively. Let $n_{1}, n_{2}, \ldots, n_{d}$ be the number of times each of the $d$ distinct values occur in $\mathbf{S}(I)$ (thus $\left.\sum_{i} n_{i}=n\right)$.

Theorem 2. There exist constants $c_{1}$ and $c_{2}$ independent of $n$ and $d$ such that

$$
\begin{gathered}
\operatorname{error}(\overline{\mathbf{S}}) \leq \sum_{i=1}^{d} \frac{c_{1} \log ^{3} n_{i}+c_{2}}{\epsilon^{2}} \\
\text { Thus error}(\overline{\mathbf{S}})=O\left(d \log ^{3} n / \epsilon^{2}\right) \text { whereas error }(\tilde{\mathbf{S}})=\Theta\left(n / \epsilon^{2}\right)
\end{gathered}
$$

The above theorem shows that constrained inference can boost accuracy, and the improvement depends on properties of the input $\mathbf{S}(I)$. In particular, if the number of distinct elements $d$ is 1 , then $\operatorname{error}(\overline{\mathbf{S}})=O\left(\log ^{3} n / \epsilon^{2}\right)$, while $\operatorname{error}(\tilde{\mathbf{S}})=\Theta\left(n / \epsilon^{2}\right)$. On the other hand, if $d=n$, then $\operatorname{error}(\overline{\mathbf{S}})=O\left(n / \epsilon^{2}\right)$ and thus both $\operatorname{error}(\overline{\mathbf{S}})$ and $\operatorname{error}(\tilde{\mathbf{S}})$ scale linearly in $n$. For many practical applications, $d_{\tilde{\mathbf{S}})} n$, which makes $\operatorname{error}(\overline{\mathbf{S}})$ significantly lower than $\operatorname{error}(\tilde{\mathbf{S}})$. In Sec. 5, experiments on real data demonstrate that the error of $\overline{\mathbf{S}}$ can be orders of magnitude lower than that of $\tilde{\mathbf{S}}$.

## 4. UNIVERSAL HISTOGRAMS

While the query sequence $\mathbf{L}$ is the conventional strategy for computing a universal histogram, this strategy has limited utility under differential privacy. While accurate for small ranges, the noise in the unit-length counts accumulates under summation, so for larger ranges, the estimates can easily become useless.

We propose an alternative query sequence that, in addition to asking for unit-length intervals, asks for interval counts of larger granularity. To ensure privacy, slightly more noise must be added to the counts. However, this strategy has the property that any range query can be answered via a linear combination of only a small number of noisy counts, and this makes it much more accurate for larger ranges.

Our alternative query sequence, denoted $\mathbf{H}$, consists of a sequence of hierarchical intervals. Conceptually, these intervals are arranged in a tree $T$. Each node $v \in T$ corresponds to an interval, and each node has $k$ children, corresponding to $k$ equally sized subintervals. The root of the tree is the interval $\left[x_{1}, x_{n}\right]$, which is recursively divided into subintervals until, at leaves of the tree, the intervals are unitlength, $\left[x_{1}\right],\left[x_{2}\right], \ldots,\left[x_{n}\right]$. For notational convenience, we define the height of the tree $\ell$ as the number of nodes, rather than edges, along the path from a leaf to the root. Thus, $\ell=\log _{k} n+1$. To transform the tree into a sequence, we arrange the interval counts in the order given by a breadthfirst traversal of the tree.


Figure 4: The tree $T$ associated with query $\mathbf{H}$ for the example in Fig. 2 for $k=2$.

Example 6. Continuing from the example in Fig 2, we describe $\mathbf{H}$ for the src domain. The intervals are arranged into a binary $(k=2)$ tree, as shown in Fig 4. The root is associated with the interval $[0 * *]$, which is evenly subdivided among its children. The unit-length intervals at the leaves are [000], [001], [010], [011]. The height of the tree is $\ell=3$.

The intervals of the tree are arranged into the query sequence $\mathbf{H}=\left\langle C_{0 * *}, C_{00 *}, C_{01 *}, C_{000}, C_{001}, C_{010}, C_{011}\right\rangle$. Evaluated on instance I from Fig. 2, the answer is $\mathbf{H}(I)=$ $\langle 14,2,12,2,0,10,2\rangle$.
To answer $\mathbf{H}$ under differential privacy, we must determine its sensitivity. As the following proposition shows, $\mathbf{H}$ has a larger sensitivity than $\mathbf{L}$.

Proposition 4. The sensitivity of $\mathbf{H}$ is $\ell$.
By Propositions 1 and 4, the following algorithm is $\epsilon$ differentially private:

$$
\tilde{\mathbf{H}}(I)=\mathbf{H}(I)+\langle\operatorname{Lap}(\ell / \epsilon)\rangle^{m}
$$

where $m$ is the length of sequence $\mathbf{H}$, equal to the number of counts in the tree.

To answer a range query using $\tilde{\mathbf{H}}$, a natural strategy is to sum the fewest number of sub-intervals such that their union equals the desired range. However, one challenge with this approach is inconsistency: in the corresponding tree of noisy answers, there may be a parent count that does not equal to the sum of its children. This can be problematic: for example, an analyst might ask one interval query and then ask for a sub-interval and receive a larger count.

We next look at how to use the arithmetic constraints between parent and child counts (denoted $\gamma_{\mathbf{H}}$ ) to derive a consistent, and more accurate, estimate $\overline{\mathbf{H}}$.

### 4.1 Constrained Inference: Computing $\overline{\mathbf{H}}$

The analyst receives $\tilde{h}=\tilde{\mathbf{H}}(I)$, the noisy output from the differentially private algorithm $\tilde{\mathbf{H}}$. We now consider the problem of finding the minimum $L_{2}$ solution: to find the $\bar{h}$ that minimizes $\|\bar{h}-\bar{h}\|_{2}$ and also satisfies the consistency constraints $\gamma_{\mathbf{H}}$.

This problem can be viewed as an instance of linear regression. The unknowns are the true counts of the unit-length intervals. Each answer in $\tilde{h}$ is a fixed linear combination of the unknowns, plus random noise. Finding $\bar{h}$ is equivalent to finding an estimate for the unit-length intervals. In fact, $\bar{h}$ is the familiar least squares solution.

Although the least squares solution can be computed via linear algebra, the hierarchical structure of this problem instance allows us to derive an intuitive closed form solution that is also more efficient to compute, requiring only linear time. Let $T$ be the tree corresponding to $\tilde{h}$; abusing notation, for node $v \in T$, we write $\tilde{h}[v]$ to refer to the interval associated with $v$.

First, we define a possibly inconsistent estimate $z[v]$ for each node $v \in T$. The consistent estimate $\bar{h}[v]$ is then described in terms of the $z[v]$ estimates. $z[v]$ is defined recursively from the leaves to the root. Let $l$ denote the height of node $v$ and $\operatorname{succ}(v)$ denote the set of $v$ 's children.

$$
z[v]=\left\{\begin{array}{l}
\tilde{h}[v], \text { if } \mathrm{v} \text { is a leaf node } \\
\frac{k^{l}-k^{l}-1}{k^{l}-1} \tilde{h}[v]+\frac{k^{l-1}-1}{k^{l}-1} \sum_{u \in \operatorname{succ}(v)} z[u], \text { o.w. }
\end{array}\right.
$$

The intuition behind $z[v]$ is that it is a weighted average of two estimates for the count at $v$; in fact, the weights are inversely proportional to the variance of the estimates.

The consistent estimate $\bar{h}$ is defined recursively from the root to the leaves. At the root $r, \bar{h}[r]$ is simply $z[r]$. As we descend the tree, if at some node $u$, we have $\bar{h}[u] \neq$ $\sum_{w \in \operatorname{succ}(u)} z[w]$, then we adjust the values of each descendant by dividing the difference $\bar{h}[u]-\sum_{w \in \operatorname{succ}(u)} z[w]$ equally among the $k$ descendants. The following theorem states that this is the minimum $L_{2}$ solution.

Theorem 3. Given the noisy sequence $\tilde{h}=\tilde{\mathbf{H}}(I)$, the unique minimum $L_{2}$ solution, $\bar{h}$, is given by the following recurrence relation. Let $u$ be $v$ 's parent:

$$
\bar{h}[v]=\left\{\begin{array}{l}
z[v], \text { if } v \text { is the root } \\
z[v]+\frac{1}{k}\left(\bar{h}[u]-\sum_{w \in \operatorname{succ}(u)} z[w]\right), \text { o.w. }
\end{array}\right.
$$

Theorem 3 shows that the overhead of computing $\overline{\mathbf{H}}$ is minimal, requiring only two linear scans of the tree: a bottom up scan to compute $z$ and then a top down scan to compute the solution $\bar{h}$ given $z$.

### 4.2 Utility Analysis: the Accuracy of $\overline{\mathbf{H}}$

We measure utility as accuracy of range queries, and we compare three strategies: $\tilde{\mathbf{L}}, \tilde{\mathbf{H}}$, and $\overline{\mathbf{H}}$. We start by comparing $\tilde{\mathbf{L}}$ and $\tilde{\mathbf{H}}$.

Given range query $q=c([x, y])$, we derive an estimate for the answer as follows. For $\tilde{\mathbf{L}}$, the estimate is simply the sum of the noisy unit-length intervals in the range: $\tilde{\mathbf{L}}_{q}=$ $\sum_{i=x}^{y} \tilde{\mathbf{L}}[i]$. The error of each count is $2 / \epsilon^{2}$, and so the error for the range is $\operatorname{error}\left(\tilde{\mathbf{L}}_{q}\right)=O\left((y-x) / \epsilon^{2}\right)$.

For $\tilde{\mathbf{H}}$, we choose the natural strategy of summing the fewest sub-intervals of $\tilde{\mathbf{H}}$. Let $r_{1}, \ldots, r_{t}$ be the roots of disjoint subtrees of $T$ such that the union of their ranges equals
$[x, y]$. Then $\tilde{\mathbf{H}}_{q}$ is defined as $\tilde{\mathbf{H}}_{q}=\sum_{i=1}^{t} \tilde{\mathbf{H}}\left[r_{i}\right]$. Each noisy count has error equal to $2 \ell^{2} / \epsilon^{2}$ (equal to the variance of the added noise) and the number of subtrees is at most $2 \ell$ (at most two per level of the tree), thus error $\left(\tilde{\mathbf{H}}_{q}\right)=O\left(\ell^{3} / \epsilon^{2}\right)$.

There is clearly a tradeoff between these two strategies. While $\tilde{\mathbf{L}}$ is accurate for small ranges, error grows linearly with the size of the range. In contrast, the error of $\tilde{\mathbf{H}}$ is poly-logarithmic in the size of the domain (recall that $\ell=$ $\Theta(\log n))$. Thus, while $\tilde{\mathbf{H}}$ is less accurate for small ranges, it is much more accurate for large ranges. If the goal of a universal histogram is to bound worst-case or total error for all range queries, then $\tilde{\mathbf{H}}$ is the preferred strategy.

We now compare $\tilde{\mathbf{H}}$ to $\overline{\mathbf{H}}$. Since $\overline{\mathbf{H}}$ is consistent, range queries can be easily computed by summing the unit-length counts. In addition to being consistent, it is also more accurate. In fact, it is in some sense optimal: among the class of strategies that (a) produce unbiased estimates for range queries and (b) derive the estimate from linear combinations of the counts in $\tilde{h}$, there is no strategy with lower mean squared error than $\overline{\mathbf{H}}$.

Theorem 4. (i) $\overline{\mathbf{H}}$ is a linear unbiased estimator, (ii) $\operatorname{error}\left(\overline{\mathbf{H}}_{q}\right) \leq \operatorname{error}\left(E_{q}\right)$ for all $q$ and for all linear unbiased estimators $E$, (iii) error $\left(\overline{\mathbf{H}}_{q}\right)=O\left(\ell^{3} / \epsilon^{2}\right)$ for all $q$, and (iv) there exists a query $q$ s.t. error $\left(\overline{\mathbf{H}}_{q}\right) \leq \frac{3}{2(\ell-1)(k-1)-k} \operatorname{error}\left(\tilde{\mathbf{H}}_{q}\right)$.

Part (iv) of the theorem shows that $\overline{\mathbf{H}}$ can more accurate than $\tilde{\mathbf{H}}$ on some range queries. For example, in a height 16 binary tree - the kind of tree used in the experiments - there is a query $q$ where $\overline{\mathbf{H}}_{q}$ is more accurate than $\tilde{\mathbf{H}}_{q}$ by a factor of $\frac{2(\ell-1)(k-1)-k}{3}=9.33$.

Furthermore, the fact that $\overline{\mathbf{H}}$ is consistent can lead to additional advantages when the domain is sparse. We propose a simple extension to $\overline{\mathbf{H}}$ : after computing $\bar{h}$, if there is a subtree rooted at $v$ such that $\bar{h}[v] \leq 0$, we simply set the count at $v$ and all children of $v$ to be zero. This is a heuristic strategy; incorporating non-negativity constraints into inference is left for future work. Nevertheless, we show in experiments, that this small change can greatly reduce error in sparse regions and can lead to $\overline{\mathbf{H}}$ being more accurate than $\tilde{\mathbf{L}}$ even at small ranges.

## 5. EXPERIMENTS

We evaluate our techniques on three real datasets (explained in detail in Appendix C): NetTrace is derived from an IP-level network trace collected at a major university; Social Network is a graph derived from friendship relations in an online social network site; Search Logs is a dataset of search query logs over time from Jan. 1, 2004 to the present. Source code for the algorithms is available at the first author's website.

### 5.1 Unattributed Histograms

The first set of experiments evaluates the accuracy of constrained inference on unattributed histograms. We compare $\overline{\mathbf{S}}$ to the baseline approach $\tilde{\mathbf{S}}$. Since $\tilde{s}=\tilde{\mathbf{S}}(I)$ is likely to be inconsistent-out-of-order, non-integral, and possibly negative-we consider a second baseline technique, denoted $\tilde{\mathbf{S}}_{\mathbf{r}}$, which enforces consistency by sorting $\tilde{s}$ and rounding each count to the nearest non-negative integer.

We evaluate the performance of these estimators on three queries from the three datasets. On NetTrace: the query


Figure 5: Error across varying datasets and $\epsilon$. Each triplet of bars represents the three estimators: $\tilde{\mathbf{S}}$ (light gray), $\tilde{\mathrm{S}}_{\mathrm{r}}$ (gray), and $\overline{\mathrm{S}}$ (black).
returns the number of internal hosts to which each external host is connected ( $\approx 65 \mathrm{~K}$ external hosts); On Social Network, the query returns the degree sequence of the graph ( $\approx 11 \mathrm{~K}$ nodes). On Search Logs, the query returns the search frequency over a 3 -month period of the top 20 K keywords; position $i$ in the answer vector is the number of times the $i^{t h}$ ranked keyword was searched.

To evaluate the utility of an estimator, we measure its squared error. Results report the average squared error over 50 random samples from the differentially-private mechanism (each sample produces a new $\tilde{s}$ ). We also show results for three settings of $\epsilon=\{1.0,0.1,0.01\}$; smaller $\epsilon$ means more privacy, hence more random noise.

Fig 5 shows the results of the experiment. Each bar represents average performance for a single combination of dataset, $\epsilon$, and estimator. The bars represent, from left-toright, $\tilde{\mathbf{S}}$ (light gray), $\tilde{\mathbf{S}}_{\mathbf{r}}$ (gray), and $\overline{\mathbf{S}}$ (black). The vertical axis is average squared error on a log-scale. The results indicate that the proposed approach reduces the error by at least an order of magnitude across all datasets and settings of $\epsilon$. Also, the difference between $\tilde{\mathbf{S}}_{\mathbf{r}}$ and $\overline{\mathbf{S}}$ suggests that the improvement is due not simply to enforcing integrality and non-negativity but from the way consistency is enforced through constrained inference (though $\overline{\mathbf{S}}$ and $\tilde{\mathbf{S}}_{\mathbf{r}}$ are comparable on Social Network at large $\epsilon$ ). Finally, the relative accuracy of $\overline{\mathbf{S}}$ improves with decreasing $\epsilon$ (more noise). Appendix C provides intuition for how $\overline{\mathbf{S}}$ reduces error.

### 5.2 Universal Histograms

We now evaluate the effectiveness of constrained inference for the more general task of computing a universal histogram and arbitrary range queries. We evaluate three techniques for supporting universal histograms $\tilde{\mathbf{L}}, \tilde{\mathbf{H}}$, and $\overline{\mathbf{H}}$. For all three approaches, we enforce integrality and non-negativity by rounding to the nearest non-negative integer. With $\overline{\mathbf{H}}$, rounding is done as part of the inference process, using the approach described in Sec 4.2.

We evaluate the accuracy over a set of range queries of varying size and location. The range sizes are $2^{i}$ for $i=$ $1, \ldots, \ell-2$ where $\ell$ is the height of the tree. For each fixed size, we select the location uniformly at random. We report the average error over 50 random samples of $\tilde{l}$ and $\tilde{h}$, and for each sample, 1000 randomly chosen ranges.

We evaluate the following histogram queries: On NetTrace, the number of connections for each external host.


Figure 6: A comparison estimators $\tilde{\mathbf{L}}$ (circles), $\tilde{\mathbf{H}}$ (diamonds), and $\overline{\mathbf{H}}$ (squares) on two real-world datasets (top NetTrace, bottom Search Logs).

This is similar to the query in Sec 5.1 except that here the association between IP address and count is retained. On Search Logs, the query reports the temporal frequency of the query term "Obama" from Jan. 1, 2004 to present. (A day is evenly divided into 16 units of time.)

Fig 6 shows the results for both datasets and varying $\epsilon$. The top row corresponds to NetTrace, the bottom to Search Logs. Within a row, each plot shows a different setting of $\epsilon \in\{1.0,0.1,0.01\}$. For all plots, the x -axis is the size of the range query, and the $y$-axis is the error, averaged over sampled counts and intervals. Both axes are in log-scale.

First, we compare $\tilde{\mathbf{L}}$ and $\tilde{\mathbf{H}}$. For unit-length ranges, $\tilde{\mathbf{L}}$ yields more accurate estimates. This is unsurprising since it is a lower sensitivity query and thus less noise is added for privacy. However, the error of $\tilde{\mathbf{L}}$ increases linearly with the size of the range. The average error of $\tilde{\mathbf{H}}$ increases slowly with the size of the range, as larger ranges typically require summing a greater number of subtrees. For ranges larger than about 2000 units, the error of $\tilde{\mathbf{L}}$ is higher than $\tilde{\mathbf{H}}$; for the largest ranges, the error of $\tilde{\mathbf{L}}$ is $4-8$ times larger than that of $\tilde{\mathbf{H}}$ (note the log-scale).

Comparing $\overline{\mathbf{H}}$ against $\tilde{\mathbf{H}}$, the error of $\overline{\mathbf{H}}$ is uniformly lower across all range sizes, settings of $\epsilon$, and datasets. The relative performance of the estimators depends on $\epsilon$. At smaller $\epsilon$, the estimates of $\overline{\mathbf{H}}$ are more accurate relative to $\tilde{\mathbf{H}}$ and $\tilde{\mathbf{L}}$. Recall that as $\epsilon$ decreases, noise increases. This suggests that the relative benefit of statistical inference increases with the uncertainty in the observed data.

Finally, the results show that $\overline{\mathbf{H}}$ can have lower error than $\tilde{\mathbf{L}}$ over small ranges, even for leaf counts. This may be surprising since for unit-length counts, the scale of the noise of $\overline{\mathbf{H}}$ is larger than that of $\overline{\mathbf{L}}$ by a factor of $\log n$. The reduction in error is because these histograms are sparse. When the histogram contains sparse regions, $\overline{\mathbf{H}}$ can effectively identify them because it has noisy observations at higher levels of the tree. In contrast, $\tilde{\mathbf{L}}$ has only the leaf counts; thus, even if a range contains no records, $\tilde{\tilde{L}}$ will assign a positive count to roughly half of the leaves in the range.

## 6. RELATED WORK

Dwork has written comprehensive reviews of differential privacy [7, 8]; below we highlight results closest to this work.

The idea of post-processing the output of a differentially private mechanism to ensure consistency was introduced in Barak et al. [1], who proposed a linear program for making a set of marginals consistent, non-negative, and integral. However, unlike the present work, the post-processing is not shown to improve accuracy.

Blum et al. [4] propose an efficient algorithm to publish synthetic data that is useful for range queries. In comparison with our hierarchical histogram, the technique of Blum et al. scales slightly better (logarithmic versus poly-logarithmic) in terms of domain size (with all else fixed). However, our hierarchical histogram achieves lower error for a fixed domain, and importantly, the error does not grow as the size of the database increases, whereas with Blum et al. it grows with $O\left(N^{2 / 3}\right)$ with $N$ being the number of records (details in Appendix E).

The present work first appeared as an arXiv preprint [13], and since then a number of related works have emerged, including additional work by the authors. The technique for unattributed histograms has been applied to accurately and efficiently estimate the degree sequences of large social networks [12]. Several techniques for histograms over hierarchical domains have been developed. Xiao et al. [22] propose an approach based on the Haar wavelet, which is conceptually similar to the $\mathbf{H}$ query in that it is based on a tree of queries where each level in the tree is an increasingly finegrained summary of the data. In fact, that technique has error equivalent to a binary $\mathbf{H}$ query, as shown by Li et al. [15], who represent both techniques as applications of the matrix mechanism, a framework for computing workloads of linear counting queries under differential privacy. We are aware of ongoing work by McSherry et al. [17] that combines hierarchical querying with statistical inference, but differs from $\mathbf{H}$ in that it is adaptive. Chan et al. [5] consider the problem of continual release of aggregate statistics over streaming private data, and propose a differentially private counter that is similar to $\mathbf{H}$, in which items are hierarchically aggregated by arrival time. The $\mathbf{H}$ and wavelet strategy are specifically designed to support range queries. Strategies for answering more general workloads of queries under differential privacy are emerging, in both the offline [11, 15] and online [20] settings.

## 7. CONCLUSIONS

Our results show that by transforming a differentiallyprivate output so that it is consistent, we can boost accuracy. Part of the innovation is devising a query set so that useful constraints hold. Then the challenge is to apply the constraints by searching for the closest consistent solution. Our query strategies for histograms have closed-form solutions for efficiently computing a consistent answer.

Our results show that conventional differential privacy approaches can add more noise than is strictly required by the privacy condition. We believe that using constraints may be an important part of finding optimal strategies for query answering under differential privacy. More discussion of the implications of our results, and possible extensions, is included in Appendix B.

## 8. ACKNOWLEDGMENTS

Hay was supported by the Air Force Research Laboratory (AFRL) and IARPA, under agreement number FA8750-

07-2-0158. Hay and Miklau were supported by NSF CNS 0627642, NSF DUE-0830876, and NSF IIS 0643681. Rastogi and Suciu were supported by NSF IIS-0627585. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusion contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the AFRL and IARPA, or the U.S. Government.

## 9. REFERENCES

[1] B. Barak, K. Chaudhuri, C. Dwork, S. Kale, F. McSherry, and K. Talwar. Privacy, accuracy, and consistency too: A holistic solution to contingency table release. In $P O D S, 2007$.
[2] R. E. Barlow, D. J. Bartholomew, J. M. Bremner, and H. D. Brunk. Statistical Inference Under Order Restrictions. John Wiley and Sons Ltd, 1972.
[3] R. E. Barlow and H. D. Brunk. The isotonic regression problem and its dual. JASA, 67(337):140-147, 1972.
[4] A. Blum, K. Ligett, and A. Roth. A learning theory approach to non-interactive database privacy. In STOC, 2008.
[5] T.-H. H. Chan, E. Shi, and D. Song. Private and continual release of statistics. In ICALP, 2010.
[6] F. R. K. Chung and L. Lu. Survey: Concentration inequalities and martingale inequalities. Internet Mathematics, 2006.
[7] C. Dwork. Differential privacy: A survey of results. In TAMC, 2008.
[8] C. Dwork. A firm foundation for private data analysis. CACM, To appear, 2010.
[9] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In TCC, 2006.
[10] A. Ghosh, T. Roughgarden, and M. Sundararajan. Universally utility-maximizing privacy mechanisms. In STOC, 2009.
[11] M. Hardt and K. Talwar. On the geometry of differential privacy. In STOC, 2010.
[12] M. Hay, C. Li, G. Miklau, and D. Jensen. Accurate estimation of the degree distribution of private networks. In ICDM, 2009.
[13] M. Hay, V. Rastogi, G. Miklau, and D. Suciu. Boosting the accuracy of differentially-private queries through consistency. CoRR, abs/0904.0942, April 2009.
[14] J. T. G. Hwang and S. D. Peddada. Confidence interval estimation subject to order restrictions. Annals of Statistics, 1994.
[15] C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor. Optimizing histogram queries under differential privacy. In $P O D S, 2010$.
[16] F. McSherry. Privacy integrated queries: An extensible platform for privacy-preserving data analysis. In SIGMOD, 2009.
[17] F. McSherry, K. Talwar, and O. Williams. Maximum likelihood data synthesis. Manuscript, 2009.
[18] M. E. J. Newman. The structure and function of complex networks. SIAM Review, 45(2):167-256, 2003.
[19] G. Owen. Game Theory. Academic Press Ltd, 1982.
[20] A. Roth and T. Roughgarden. Interactive privacy via the median mechanism. In STOC, 2010.
[21] S. D. Silvey. Statistical Inference. Chapman-Hall, 1975.
[22] X. Xiao, G. Wang, and J. Gehrke. Differential privacy via wavelet transforms. In $I C D E, 2010$.

## APPENDIX

## A. NOTATIONAL CONVENTIONS

The table below summarizes notational conventions used in the paper.

| $\mathbf{Q}$ | Sequence of counting queries |
| :--- | :--- |
| $\mathbf{L}$ | Unit-Length query sequence |
| $\mathbf{H}$ | Hierarchical query sequence |
| $\mathbf{S}$ | Sorted query sequence |
| $\gamma_{\mathbf{Q}}, \tilde{\mathbf{L}}, \tilde{\mathbf{H}}, \tilde{\mathbf{S}}$ | Constraint set for query $\mathbf{Q}$ |
| $\tilde{\mathbf{Q}}, \mathbf{S}$ | Randomized query sequence |
| $I$ | Randomized query sequence, |
| $\mathbf{L}(I), \mathbf{H}(I), \mathbf{S}(I)$ | returning minimum L2 solution |
| $\tilde{l}=\tilde{\mathbf{L}}(I), \tilde{h}=\tilde{\mathbf{H}}(I), \tilde{s}=\tilde{\mathbf{S}}(I)$ | Private database instance |
| $\bar{h}=\overline{\mathbf{H}}(I), \bar{s}=\overline{\mathbf{S}}(I)$ | Output sequence (truth) |
|  | Output sequence (noisy) |

## B. DISCUSSION OF MAIN RESULTS

Here we provide a supplementary discussion of the results, review the insights gained, and discuss future directions.

Unattributed histograms. The choice of the sorted query $\mathbf{S}$, instead of $\mathbf{L}$, is an unqualified benefit, because we gain from the inequality constraints on the output, while the sensitivity of $\mathbf{S}$ is no greater than that of $\mathbf{L}$. Among other applications, this allows for extremely accurate estimation of degree sequences of a graph, improving error by an order of magnitude on the baseline technique. The accuracy of the estimate depends on the input sequence. It works best for sequences with duplicate counts, which matches well the degree sequences of social networks encountered in practice.

Future work specifically oriented towards degree sequence estimation could include a constraint enforcing that the output sequence is graphical, i.e. the degree sequence of some graph.

Universal histograms. The choice of the hierarchical counting query $\mathbf{H}$, instead of $\mathbf{L}$, offers a trade off because the sensitivity of $\mathbf{H}$ is greater than that of $\mathbf{L}$. It is interesting that for some data sets and privacy levels, the effect of the $\mathbf{H}$ constraints outweighs the increased noise that must be added. In other cases, the algorithms based on $\mathbf{H}$ provide greater accuracy for all but the smallest ranges. We note that in many practical settings, domains are large and sparse. The sparsity implies that no differentially private technique can yield meaningful answers for unit-length queries because the noise necessary for privacy will drown out the signal. So while $\tilde{\mathbf{L}}$ sometimes has higher accuracy for small range queries, this may not have practical relevance since the relative error of the answers renders them useless.

In future work we hope to extend the technique for universal histograms to multi-dimensional range queries, and to investigate optimizations such as higher branching factors.

Across both histogram tasks, our results clearly show that it is possible to achieve greater accuracy without sacrificing privacy. The existence of our improved estimators $\overline{\mathbf{S}}$ and $\overline{\mathbf{H}}$ show that there is another differentially private noise distribution that is more accurate than independent Laplace noise. This does not contradict existing results because the original differential privacy work showed only that calibrating Laplace noise to the sensitivity of a query was sufficient for privacy, not that it was necessary. Only recently has the optimality of this construction been studied (and proven only for single queries) [10]. Finding the optimal strategy for answering a set of queries under differential privacy is an important direction for future work, especially in light of emerging private query interfaces [16].

A natural goal is to describe directly the improved noise distributions implied by $\overline{\mathbf{S}}$ and $\overline{\mathbf{H}}$, and build a privacy mechanism that samples from it. This could, in theory, avoid the inference step altogether. But it is seems quite difficult to discover, describe, and sample these improved noise distributions, which will be highly dependent on a particular query of interest. Our approach suggests that constraints and constrained inference can be an effective path to discovering new, more accurate noise distributions that satisfy differential privacy. As a practical matter, our approach does not necessarily burden the analyst with the constrained inference process because the server can implement the postprocessing step. In that case it would appear to the analyst as if the server was sampling from the improved distribution.

While our focus has been on histogram queries, the techniques are probably not limited to histograms and could have broader impact. However, a general formulation may be challenging to develop. There is a subtle relationship between constraints and sensitivity: reformulating a query so that it becomes highly constrained may similarly increase its sensitivity. A challenge is finding queries such as $\mathbf{S}$ and $\mathbf{H}$ that have useful constraints but remain low sensitivity. Another challenge is the computational efficiency of constrained inference, which is posed here as a constrained optimization problem with a quadratic objective function. The complexity of solving this problem will depend on the nature of the constraints and is NP-Hard in general. Our analysis shows that the constraint sets of $\mathbf{S}$ and $\mathbf{H}$ admit closed-form solutions that are efficient to compute.

## C. ADDITIONAL EXPERIMENTS

This section provides detailed descriptions of the datasets, and additional results for unattributed histograms.
NetTrace is derived from an IP-level network trace collected at a major university. The trace monitors traffic at the gateway between internal IP addresses and external IP addresses. From this data, we derived a bipartite connection graph where the nodes are hosts, labeled by their IP address, and an edge connotes the transmission of at least one data packet. Here, differential privacy ensures that individual connections remain private.

Social Network is a graph derived from friendship relations on an online social network site. The graph is limited to a population of roughly 11,000 students from a single university. Differential privacy implies that friendships will not be disclosed. The size of the graph (number of students) is assumed to be public knowledge. ${ }^{1}$

[^0]

Figure 7: On NetTrace, $\mathbf{S}(I)$ (solid gray), the average error of $\overline{\mathbf{S}}$ (solid black) and $\tilde{\mathbf{S}}$ (dotted gray), for $\epsilon=1.0$.

Search Logs is a dataset of search query logs over time from Jan. 1, 2004 to the present. For privacy reasons, it is difficult to obtain such data. Our dataset is derived from a search engine interface that publishes summary statistics for specified query terms. We combined these summary statistics with a second dataset, which contains actual search query logs but for a much shorter time period, to produce a synthetic data set. In the experiments, ground truth refers to this synthetic dataset. Differential privacy guarantees that the output will prevent the association of an individual entity (user, host) with a particular search term.

Unattributed histograms. Figure 7 provides some intuition for how inference is able to reduce error. Shown is a portion of the unattributed histogram of NetTrace: the sequence is sorted in descending order along the x -axis and the y -axis indicates the count. The solid gray line corresponds to ground truth: a long horizontal stretch indicates a subsequence of uniform counts and a vertical drop indicates a decrease in count. The graphic shows only the middle portion of the unattributed histogram-some very large and very small counts are omitted to improve legibility. The solid black lines indicate the error of $\overline{\mathbf{S}}$ averaged over 200 random samples of $\tilde{\mathbf{S}}$ (with $\epsilon=1.0$ ); the dotted gray lines indicate the expected error of $\tilde{\mathbf{S}}$.

The inset graph on the left reveals larger error at the beginning of the sequence, when each count occurs once or only a few times. However, as the counts become more concentrated (longer subsequences of uniform count), the error diminishes, as shown in the right inset. Some error remains around the points in the sequence where the count changes, but the error is reduced to zero for positions in the middle of uniform subsequences.

Figure 7 illustrates that our approach reduces or eliminates noise in precisely the parts of the sequence where the noise is unnecessary for privacy. Changing a tuple in the database cannot change a count in the middle of a uniform
of students can be estimated privately within $\pm 1 / \epsilon$ in expectation. Our techniques can be applied directly to either the true count or a noisy estimate.
subsequence, only at the end points. These experimental results also align with Theorem 2, which states that the error of $\overline{\mathbf{S}}$ is a function of the number of distinct counts in the sequence. In fact, the experimental results suggest that the theorem also holds locally for subsequences with a small number of distinct counts. This is an important result since the typical degree sequences that arise in real data, such as the power-law distribution, contain very large uniform subsequences.

## D. PROOFS

Proof of Proposition 2. For any output $\bar{q}$, then let $\mathcal{S}(\bar{q})$ denote the set of noisy answers such that if $\tilde{q} \in \mathcal{S}(\bar{q})$ then the minimum $L_{2}$ solution given $\tilde{q}$ and $\gamma_{\mathbf{Q}}$ is $\bar{q}$. For any $I$ and $I^{\prime} \in \operatorname{nbrs}(I)$, the following shows that $\overline{\mathbf{Q}}$ is $\epsilon$ differentially private:

$$
\begin{aligned}
\operatorname{Pr}[\overline{\mathbf{Q}}(I)=\bar{q}] & =\operatorname{Pr}[\tilde{\mathbf{Q}}(I) \in \mathcal{S}(\bar{q})] \\
& \leq \exp (\epsilon) \operatorname{Pr}\left[\tilde{\mathbf{Q}}\left(I^{\prime}\right) \in \mathcal{S}(\bar{q})\right] \\
& =\exp (\epsilon) \operatorname{Pr}\left[\overline{\mathbf{Q}}\left(I^{\prime}\right)=\bar{q}\right]
\end{aligned}
$$

where the inequality is because $\tilde{\mathbf{Q}}$ is $\epsilon$-differentially private.

Proof of Proposition 3. Given a database $I$, suppose we add a record to it to obtain $I^{\prime}$. The added record affects one count in $\mathbf{L}$, i.e., there is exactly one $i$ such that $\mathbf{L}(I)[i]=$ $x$ and $\mathbf{L}\left(I^{\prime}\right)[i]=x+1$, and all other counts are the same. The added record affects $\mathbf{S}$ as follows. Let $j$ be the largest index such that $\mathbf{S}(I)[j]=x$, then the added record increases the count at $j$ by one: $\mathbf{S}\left(I^{\prime}\right)[j]=x+1$. Notice that this change does not affect the sort order-i.e., in $\mathbf{S}\left(I^{\prime}\right)$, the $j^{\text {th }}$ value remains in sorted order: $\mathbf{S}\left(I^{\prime}\right)[j-1] \leq x, \mathbf{S}\left(I^{\prime}\right)[j]=x+1$, and $\mathbf{S}\left(I^{\prime}\right)[j+1] \geq x+1$. All other counts in $\mathbf{S}$ are the same, and thus the $L_{1}$ distance between $\mathbf{S}(I)$ and $\mathbf{S}\left(I^{\prime}\right)$ is 1 .

Proof of Proposition 4. If a tuple is added or removed from the relation, this affects the count for every range that includes it. There are exactly $\ell$ ranges that include a given tuple: the range of a single leaf and the ranges of
the nodes along the path from that leaf to the root. Therefore, adding/removing a tuple changes exactly $\ell$ counts each by exactly 1 . Thus, the sensitivity is equal to $\ell$, the height of the tree.

## D. 1 Proof of Theorem 1

We first restate the theorem below. Recall that $\tilde{s}[i, j]$ denotes the subsequence of $j-i+1$ elements: $\langle\tilde{s}[i], \tilde{s}[i+1]$, $\ldots, \tilde{s}[j]\rangle$. Let $\tilde{M}[i, j]$ record the mean of these elements, i.e. $\tilde{M}[i, j]=\sum_{k=i}^{j} \tilde{s}[k] /(j-i+1)$.

Theorem 1. Denote $L_{k}=\min _{j \in[k, n]} \max _{i \in[1, j]} \tilde{M}[i, j]$ and $U_{k}=\max _{i \in[1, k]} \min _{j \in[i, n]} \tilde{M}[i, j]$. The minimum $L_{2}$ solution $\bar{s}$, is unique and given by: $\bar{s}[k]=L_{k}=U_{k}$.

Proof. In the proof, we abbreviate the notation and implicitly assume that the range of $i$ is $[1, n]$ or $[1, j]$ when $j$ is specified. Similarly, the range of $j$ is $[1, n]$ or $[i, n]$ when $i$ is specified.

We start with the easy part, showing that $U_{k} \leq L_{k}$. Define an $n \times n$ matrix $A^{k}$ as follows:

$$
A_{i j}^{k}= \begin{cases}\tilde{M}[i, j] & \text { if } \mathrm{i} \leq \mathrm{j} \\ \infty & \text { if } \mathrm{j}<\mathrm{i} \leq \mathrm{k} \\ -\infty & \text { otherwise }\end{cases}
$$

Then $\min _{j} \max _{i} A_{i j}^{k}=L_{k}$ and $\max _{i} \min _{j} A_{i j}^{k}=U_{k}$. In any matrix $A^{k}, \max _{i} \min _{j} A_{i j}^{k} \leq \min _{j} \max _{i} A_{i j}^{k}$ : this is a simple fact that can be checked directly, or see [19], hence $U_{k} \leq L_{k}$.

We show next that if $\bar{s}$ is the minimum $L_{2}$ solution, then $L_{k} \leq \bar{s}[k] \leq U_{k}$. If we show this, then the proof of the theorem is completed, as then we will then have $\bar{s}[k]=L_{k}=$ $U_{k}$. The proof relies on the following lemma.

Lemma 1. Let $\bar{s}$ be the minimum $L_{2}$ solution. Then (i) $\bar{s}[1] \leq U_{1}$, (ii) $\bar{s}[n] \geq L_{n}$, (iii) for all $k$, $\min \left(\bar{s}[k+1], \max _{i} \tilde{M}[i, k]\right)$ $\bar{s}[k] \leq \max \left(\bar{s}[k-1], \min _{j} \tilde{M}[k, j]\right)$.

The proof of the lemma appears below, but now we use it to complete the proof of Theorem 1. First, we show that $\bar{s}[k] \leq U_{k}$ using induction on $k$. The base case is $k=1$ and it is stated in the lemma, part (i). For the inductive step, assume $\bar{s}[k-1] \leq U_{k-1}$. From (iii), we have that

$$
\begin{aligned}
\bar{s}[k] & \leq \max \left(\bar{s}[k-1], \min _{j} \tilde{M}[k, j]\right) \\
& \leq \max \left(U_{k-1}, \min _{j} \tilde{M}[k, j]\right)=U_{k}
\end{aligned}
$$

The last step follows from the definition of $U_{k}$. A similar induction argument shows that $\bar{s}[k] \geq L_{k}$, except the order is reversed: the base case is $k=n$ and the inductive step assumes $\bar{s}[k+1] \geq L_{k+1}$.

The only remaining step is to prove the lemma.
Proof of Lemma 1. For (i), it is sufficient to prove that $\bar{s}[1] \leq \tilde{M}[1, j]$ for all $j \in[1, n]$. Assume the contrary. Thus there exists a $j$ such that for $\bar{s}[1]>\tilde{M}[1, j]$. Let $\delta=\bar{s}[1]-$ $\tilde{M}[1, j]$. Thus $\delta>0$. Further, for all $i$, denote $\delta_{i}=\bar{s}[i]-\bar{s}[1]$. Consider the sequence $\bar{s}^{\prime}$ defined as follows:

$$
\bar{s}^{\prime}[i]= \begin{cases}\bar{s}[i]-\delta & \text { if } i \leq j \\ \bar{s}[i] & \text { otherwise }\end{cases}
$$

It is obvious to see that since $\bar{s}$ is a sorted sequence, so is $\bar{s}^{\prime}$.

We now claim that $\left\|\bar{s}^{\prime}-\tilde{s}\right\|_{2}<\|\bar{s}-\tilde{s}\|_{2}$. For this note that since the sequence $\bar{s}^{\prime}[j+1, n]$ is identical to the sequence $\bar{s}[j+$ $1, n]$, it is sufficient to prove $\left\|\bar{s}^{\prime}[1, j]-\tilde{s}[1, j]\right\|_{2}<\| \bar{s}[1, j]-$ $\tilde{s}[1, j] \|_{2}$. To prove that, note that $\|\bar{s}[1, j]-\tilde{s}[1, j]\|_{2}$ can be expanded as

$$
\begin{aligned}
\|\bar{s}[1, j]-\tilde{s}[1, j]\|_{2} & =\sum_{i=1}^{j}(\bar{s}[i]-\tilde{s}[i])^{2}=\sum_{i=1}^{j}\left(\bar{s}[1]+\delta_{i}-\tilde{s}[i]\right)^{2} \\
& =\sum_{i=1}^{j}\left(\tilde{M}[1, j]+\delta+\delta_{i}-\tilde{s}[i]\right)^{2}
\end{aligned}
$$

Suppose for a moment that we fix $\tilde{M}[1, j]$ and $\delta_{i}$ 's, and treat $\|\bar{s}[1, j]-\tilde{s}[1, j]\|_{2}$ as a function $f$ over $\delta$. The derivative of $f(\delta)$ is:

$$
\begin{aligned}
f^{\prime}(\delta) & =2 \sum_{i=1}^{j}\left(\tilde{M}[1, j]+\delta+\delta_{i}-\tilde{s}[i]\right) \\
& =2\left(j \tilde{M}[1, j]-\sum_{i=1}^{j} \tilde{s}[i]\right)+2 j \delta+2 \sum_{i=1}^{j} \delta_{i} \\
& =2 j \delta+2 \sum_{i=1}^{j} \delta_{i}
\end{aligned}
$$

Since $\delta_{i} \geq 0$ for all $i$, then the derivative is strictly greater than zero for any $\delta>0$, which implies that $f$ is a strictly increasing function of $\delta$ and has a minimum at $\delta=0$. Therefore, $\|\bar{s}[1, j]-\tilde{s}[1, j]\|_{2}=f(\delta)>f(0)=\left\|\bar{s}^{\prime}[1, j]-\tilde{s}[1, j]\right\|_{2}$. This is a contradiction since it was assumed that $\bar{s}$ was the minimum solution. This completes the proof for (i).

For (ii), the proof of $\bar{s}[n] \geq \max _{i} \tilde{M}[i, n]$ follows from a similar argument: if $\bar{s}[n]<\tilde{M}[i, n]$ for some $i$, define $\delta=\tilde{M}[i, n]-\bar{s}[n]$ and the sequence $\bar{s}^{\prime}$ with elements $\bar{s}^{\prime}[j]=$ $s[j]+\delta$ for $j \geq i$. Then $\bar{s}^{\prime}$ can be shown to be a strictly better solution than $\bar{s}$, proving (ii).

For the proof of (iii), we first show that $\bar{s}[k] \leq \max (\bar{s}[k-$ $1], \min _{j} \tilde{M}[k, j]$ ). Assume the contrary, i.e. there exists a $k$ such that $\bar{s}[k]>\bar{s}[k-1]$ and $\bar{s}[k]>\min _{j} \tilde{M}[k, j]$. In other words, we assume there exists a $k$ and $j$ such that $\bar{s}[k]>\bar{s}[k-1]$ and $\bar{s}[k]>\tilde{M}[k, j]$. Denote $\delta=\bar{s}[k]-$ $\max (\bar{s}[k-1], \tilde{M}[k, j])$. By our assumption above, $\delta>0$. Define the sequence

$$
\bar{s}^{\prime}[i]= \begin{cases}\bar{s}[i]-\delta & \text { if } k \leq i \leq j \\ \bar{s}[i] & \text { otherwise }\end{cases}
$$

Note that by construction, $\bar{s}^{\prime}[k]=\bar{s}[k]-\delta=\bar{s}[k]-(\bar{s}[k]-$ $\max (\bar{s}[k-1], \tilde{M}[k, j]))=\max (\bar{s}[k-1], \tilde{M}[k, j])$. It is easy to see that $\bar{s}^{\prime}$ is sorted (indeed the only inversion in the sort order could have occurred if $\bar{s}^{\prime}[k-1]>\bar{s}^{\prime}[k]$, but doesn't as $\left.\bar{s}^{\prime}[k-1]=\bar{s}[k-1] \leq \max (\bar{s}[k-1], \tilde{M}[k, j])=\bar{s}^{\prime}[k]\right)$.

Now a similar argument as in the proof of (i) for the sequence $\tilde{s}[k, j]$, yields that the error $\left\|\bar{s}^{\prime}[k, j]-\tilde{s}[k, j]\right\|_{2}<$ $\|\bar{s}[k, j]-\tilde{s}[k, j]\|_{2}$. Thus $\left\|\bar{s}^{\prime}-\tilde{s}\right\|_{2}<\left\|\bar{s}^{\prime}-\tilde{s}\right\|_{2}$ and $\bar{s}^{\prime}$ is a strictly better solution than $\bar{s}$. This yields a contradiction as $\bar{s}$ is the minimum $L_{2}$ solution. Hence $\bar{s}[k] \leq$ $\max \left(\bar{s}[k-1], \min _{j} \tilde{M}[k, j]\right)$.

A similar argument in the the reverse direction shows that $\bar{s}[k] \geq \min \left(\bar{s}_{k+1}, \max _{i} \tilde{M}[i, k]\right)$ completing the proof of (iii).

## D. 2 Proof of Theorem 2

We first restate the theorem below. Denote $n$ and $d$ as the number of values and the number of distinct values in $\mathbf{S}(I)$ respectively. Let $n_{1}, n_{2}, \ldots, n_{d}$ be the number of times each of the $d$ distinct values occur in $\mathbf{S}(I)$ (thus $\sum_{i} n_{i}=n$ ).

Theorem 2. There exist constants $c_{1}$ and $c_{2}$ independent of $n$ and $d$ such that

$$
\operatorname{error}(\overline{\mathbf{S}}) \leq \sum_{i=1}^{d} \frac{c_{1} \log ^{3} n_{i}+c_{2}}{\epsilon^{2}}
$$

Thus error $(\overline{\mathbf{S}})=O\left(d \log ^{3} n / \epsilon^{2}\right)$ whereas error $(\tilde{\mathbf{S}})=\Theta\left(n / \epsilon^{2}\right)$.
Before showing the proof, we prove the following lemma.
Lemma 2. Let $s=\mathbf{S}(I)$ be the input sequence. Call a translation of $s$ the operation of subtracting from each element of $s$ a fixed amount $\delta$. Then error $(\overline{\mathbf{S}}[i])$ is invariant under translation for all $i$.

Proof. Denote $\operatorname{Pr}(\bar{s} \mid s)(\operatorname{Pr}(\tilde{s} \mid s))$ the probability that $\bar{s}(\tilde{s})$ is output on the input sequence $s$. Denote $s^{\prime}, \bar{s}^{\prime}$, and $\tilde{s}^{\prime}$ the sequence obtained by translating $s, \bar{s}$, and $\tilde{s}$ by $\delta$, respectively.

First observe that $\operatorname{Pr}(\tilde{s} \mid s)=\operatorname{Pr}\left(\tilde{s}^{\prime} \mid s^{\prime}\right)$ as $\tilde{s}$ and $\tilde{s}^{\prime}$ are obtained by adding the same Laplacian noise to $s$ and $s^{\prime}$, respectively. Using Theorem 1 (since all $U_{k}$ 's and $L_{k}$ 's shift by $\delta$ on translating $\tilde{s}$ by delta), we get that if $\bar{s}$ is the minimum $L_{2}$ solution given $\tilde{s}$, then $\bar{s}^{\prime}$ is the minimum $L_{2}$ solution given $\tilde{s}^{\prime}$. Thus, $\operatorname{Pr}(\bar{s} \mid s)=\operatorname{Pr}\left(\bar{s}^{\prime} \mid s^{\prime}\right)$ for all sequences $\bar{s}$. Further, since $\bar{s}[i]$ and $\bar{s}^{\prime}[i]$ yield the same $L_{2}$ error with $s[i]$ and $s^{\prime}[i]$ respectively, we get that the expected $\operatorname{error}(\overline{\mathbf{S}}[i])$ is same for both inputs $s$ and $s^{\prime}$.

Lemma 3. Let $X$ be any positive random variable that is bounded ( $\lim _{x \rightarrow \infty} x \operatorname{Pr}(X>x)$ exists). Then

$$
\mathbb{E}(X) \leq \int_{0}^{\infty} \operatorname{Pr}(X>x) d x
$$

Proof. The proof follows from the following chain of equalities.

$$
\begin{aligned}
& \mathbb{E}(X)=\int_{0}^{\infty} x \frac{\partial}{\partial x}(\operatorname{Pr}(X \leq x)) \\
= & -\int_{0}^{\infty} x \frac{\partial}{\partial x}(\operatorname{Pr}(X>x)) \\
= & -[x \operatorname{Pr}(X>x)]_{0}^{\infty}+\int_{0}^{\infty}(\operatorname{Pr}(X \leq x)-1) d x \quad \text { (by parts) } \\
= & -\lim _{x \rightarrow \infty} x \operatorname{Pr}(X>x)+\int_{0}^{\infty} \operatorname{Pr}(X>x) d x \\
\leq & \int_{0}^{\infty} \operatorname{Pr}(X>x) d x
\end{aligned}
$$

Here the last equality follows as $X$ is bounded and therefore the limit exists and is positive. This completes the proof.

We next state a theorem that was shown in [6]
Theorem 5 (Theorem 3.4 [6]). Suppose that $X_{1}, X_{2}$, $\ldots, X_{n}$ are independent random variables satisfying $X_{i} \leq$ $\mathbb{E}\left(X_{i}\right)+M$, for $1 \leq i \leq n$. We consider the sum $X=$
$\sum_{i=1}^{n} X_{i}$ with expectation $\mathbb{E}(X)=\sum_{i=1}^{n} E\left(X_{i}\right)$ and $\operatorname{Var}(X)=$ $\sum_{i=1}^{n=1} \operatorname{Var}\left(X_{i}\right)$. Then, we have

$$
\operatorname{Pr}(X \geq \mathbb{E}(X)+\lambda) \leq e^{\frac{-\lambda^{2}}{2(\operatorname{Var}(X)+M \lambda / 3)}}
$$

For a random variable $X$, denote $\mathbb{I}_{X}$ the indicator function that $X \geq 0$ (thus $\mathbb{I}_{X}=1$ if $X \geq 0$ and 0 otherwise). Using Theorem 5 , we prove the following lemma.

Lemma 4. Suppose $i, j$ are indices such that for all $k \in$ $[i, j], s[k] \leq 0$. Then there exists a constant $c$ such that for all $\tau \geq 1$ the following holds.
$\operatorname{Pr}\left(\tilde{M}[i, j]^{2} \mathbb{I}_{\tilde{M}[i, j]} \geq c\left(\frac{\log ^{2}((j-i+1) \tau)}{(j-i+1) \epsilon^{2}}\right)\right) \leq \frac{1}{(j-i+1)^{2} \tau^{2}}$
Proof. We apply Theorem 5 on $\tilde{s}[k]$ for $k \in[i, j]$. First note that $\mathbb{E}(\tilde{s}[k])=s[k] \leq 0$. Further $\operatorname{Var}(\tilde{s}[k])=\frac{2}{\epsilon^{2}}$ as $\tilde{s}[k]$ is obtained by adding Laplace noise to $s[k]$ which $\epsilon^{\epsilon^{2}}$ has this variance. We also know that $\tilde{s}[k] \geq M+s[k]$ happens with probability at most $e^{-\epsilon M} / 2$.

For simplicity, call $n$ to be $j-i+1$. Denoting $X=$ $\sum_{k \in[i, j]} \tilde{s}[k]$, we see that $\mathbb{E}(X) \leq 0$ and $\operatorname{Var}(X)=\frac{2 n}{\epsilon^{2}}$. Further, set $M=3 \log (n \tau) / \epsilon$. Denote $B$ the event that for some $k, \tilde{s}[k] \geq M+s[k]$. Thus $\operatorname{Pr}(B) \leq n e^{-\epsilon M} / 2 \leq \frac{1}{2 n^{2} \tau^{3}}$. If $B$ does not happen, we know that $\tilde{s}[k] \leq M+s[k]$ for all $k \in[i, j]$. Thus we can then apply Theorem 5 to get:

$$
\begin{aligned}
\operatorname{Pr}(X \geq \mathbb{E}(X)+\lambda) & \leq e^{\frac{-\lambda^{2}}{2\left(2 n / \epsilon^{2}+\lambda \log (n \tau) / \epsilon\right)}}+\operatorname{Pr}(B) \\
& =e^{\frac{-\lambda^{2}}{2\left(2 n / \epsilon^{2}+\lambda \log (n \tau) / \epsilon\right)}}+\frac{1}{2 n^{2} \tau^{3}}
\end{aligned}
$$

Setting $\lambda=\frac{8}{\epsilon} \sqrt{n} \log (n \tau)$ gives us that

$$
\operatorname{Pr}\left(X \geq \mathbb{E}(X)+\frac{8}{\epsilon} \sqrt{n} \log (n \tau)\right) \leq \frac{1}{n^{2} \tau^{2}}
$$

Since $\mathbb{E}(X) \leq 0$, we get

$$
\operatorname{Pr}\left(X \geq \frac{8}{\epsilon} \sqrt{n} \log (n \tau)\right) \leq \frac{1}{n^{2} \tau^{2}}
$$

Also we observe that $\tilde{M}[i, j]=X / n$, which yields

$$
\operatorname{Pr}\left(\tilde{M}[i, j] \geq \frac{8 \log (n \tau)}{\sqrt{n} \epsilon}\right) \leq \frac{1}{n^{2} \tau^{2}}
$$

Finally, observe that $\tilde{M}[i, j] \leq c$ implies that $\tilde{M}[i, j]^{2} \mathbb{I}_{\tilde{M}[i, j]} \leq$ $c^{2}$. Thus we get

$$
\operatorname{Pr}\left(\tilde{M}[i, j]^{2} \mathbb{I}_{m m[i, j]} \geq \frac{64 \log ^{2}(n \tau)}{n \epsilon^{2}}\right) \leq \frac{1}{n^{2} \tau^{2}}
$$

Putting $n=j-i+1$ and using $c=64$ gives us the required result.

Now we can give the proof of Theorem 2.
Proof of Theorem 2. The proof of $\operatorname{error}(\tilde{\mathbf{S}})=\Theta\left(n / \epsilon^{2}\right)$ is obvious since:

$$
\operatorname{error}(\tilde{\mathbf{S}})=\sum_{k=1}^{n} \operatorname{error}(\tilde{s}[i])=n\left(\frac{2}{\epsilon^{2}}\right)
$$

In the rest of the proof, we shall show bound $\operatorname{error}(\overline{\mathbf{S}})$. Let $s=S(I)$ be the input sequence. We know that $s$ consists of $d$ distinct elements. Denote $s_{r}$ as the $r^{t h}$ distinct element of $s$. Also denote $\left[l_{r}, u_{r}\right]$ as the set of indices corresponding
to $s_{r}$, i.e. $\forall_{i \in\left[l_{r}, u_{r}\right]} s[i]=s_{r}$ and $\forall_{i \notin\left[l_{r}, u_{r}\right]} s[i] \neq s_{r}$. Let $M[i, j]$ record the mean of elements in $s[i, j]$, i.e. $M[i, j]=$ $\sum_{k=i}^{j} s[k] /(i-j+1)$.

To bound error $(\overline{\mathbf{S}})$, we shall bound $\operatorname{error}(\overline{\mathbf{S}}[i])$ separately for each $i$. To bound $\operatorname{error}(\overline{\mathbf{S}}[i])$, we can assume W.L.O.G that $s[i]$ is 0 . This is because if $s[i] \neq 0$, then we can translate the sequence $s$ by $s[i]$. As shown in Lemma 2 this preserves error $(\overline{\mathbf{S}}[i])$, while making $s[i]=0$.

Let $k \in\left[l_{r}, u_{r}\right]$ be any index for the $r^{t h}$ distinct element of $s$. By definition, $\operatorname{error}(\overline{\mathbf{S}}[k])=\mathbb{E}(\bar{s}[k]-s[k])^{2}=\mathbb{E}\left(\bar{s}[k]^{2}\right)($ as we can assume W.L.O.G $s[k]=0$ ). From Theorem 1, we know that $\bar{s}[k]=U_{k}$. Thus $\operatorname{error}(\overline{\mathbf{S}}[k])=\mathbb{E}\left(U_{k}^{2}\right)$. Here we treat $U_{k}=\max _{i \leq k} \min _{j} \tilde{M}[i, j]$ as a random variable. Now by definition of $\mathbb{E}$, we have

$$
\mathbb{E}\left(U_{k}^{2}\right)=\mathbb{E}\left(U_{k}^{2} \mathbb{I}_{U_{k}}\right)+\mathbb{E}\left(U_{k}^{2}\left(1-\mathbb{I}_{U_{k}}\right)\right)=A+B(\text { say })
$$

We shall bound $A$ and $B$ separately. For bounding $A$, denote $\mathcal{U}_{k}=\max _{i \leq k} \tilde{M}\left[i, u_{r}\right]$. It is apparent that $\mathcal{U}_{k} \geq U_{k}$ and thus $\mathcal{U}_{k}^{2} \mathbb{I}_{\mathcal{U}_{k}} \geq U_{k}^{2} \mathbb{I}_{U_{k}}$. To bound $A$, we observe that

$$
A=\mathbb{E}\left(U_{k}^{2} \mathbb{I}_{U_{k}}\right) \leq \mathbb{E}\left(\mathcal{U}_{k}^{2} \mathbb{I}_{\mathcal{U}_{k}}\right)
$$

Further, since $\mathcal{U}_{k}=\max _{i \leq k} \tilde{M}\left[i, u_{r}\right]$, we know that $\mathcal{U}_{k}^{2} \mathbb{I}_{\mathcal{U}_{k}}=$ $\max _{i \leq k} \tilde{M}\left[i, u_{r}\right]^{2} \mathbb{I}_{\tilde{M}\left[i, u_{r}\right]}$. Thus we can write:

$$
A \leq \mathbb{E}\left(\mathcal{U}_{k}^{2} \mathbb{I}_{\mathcal{U}_{k}}\right)=\mathbb{E}\left(\max _{i \leq k} \tilde{M}\left[i, u_{r}\right]^{2} \mathbb{I}_{\tilde{M}\left[i, u_{r}\right]}\right)
$$

Let $\tau>1$ be any number and $c$ be the constant used in Lemma 4. Let us denote $e_{i}$ the event that:

$$
\tilde{M}\left[i, u_{r}\right]^{2} \mathbb{I}_{\tilde{M}\left[i, u_{r}\right]} \geq c\left(\frac{\log ^{2}\left(\left(u_{r}-i+1\right) \tau\right)}{\left(u_{r}-i+1\right) \epsilon^{2}}\right)
$$

We can apply lemma 4 to compute the probability of $e_{i}$ as $s[j] \leq 0$ for all $j \leq u_{r}$ (as we assumed W.L.O.G $s[k]=0$ ). Thus we get $\operatorname{Pr}\left(e_{i}\right) \leq \frac{1}{\left(u_{r}-i+1\right)^{2} \tau^{2}}$.

Define $e=\vee_{i=1}^{u_{r}} e_{i}$. Then $\operatorname{Pr}(e) \leq \sum_{i=1}^{u_{r}} \operatorname{Pr}\left(e_{i}\right)=2 / \tau^{2}$ (as $\sum_{i=1}^{u_{r}} 1 / i^{2} \leq 2$ ). If the event $e$ does not happen, then it is easy to see that

$$
\begin{aligned}
\mathcal{U}_{k}^{2} \mathbb{I}_{\mathcal{U}_{k}} & =\max _{i \leq k} \tilde{M}\left[i, u_{r}\right]^{2} \mathbb{I}_{\tilde{M}\left[i, u_{r}\right]} \\
& \leq c\left(\frac{\log ^{2}\left(\left(u_{r}-k+1\right) \tau\right)}{\left(u_{r}-k+1\right) \epsilon^{2}}\right)
\end{aligned}
$$

Thus with at least probability $1-2 / \tau^{2}$ (which is $\operatorname{Pr}(\neg e)$ ), we get $\mathcal{U}_{k}^{2} \mathbb{I}_{\mathcal{U}_{k}}$ is bounded as above. This yields that there exist constants $c_{1}$ and $c_{2}$ such that $\mathbb{E}\left(\mathcal{U}_{k}^{2} \mathbb{I}_{\mathcal{U}_{k}}\right) \leq \frac{c_{1} \log ^{2}\left(u_{r}-k+1\right)+c_{2}}{\left(u_{r}-k+1\right) \epsilon^{2}}$. The proof is by the application of Lemma 3 (as $\mathcal{U}_{k}$ is bounded) and a simple integration over $\tau$ ranging from 1 to $\infty$. Finally we get that $A \leq \mathbb{E}\left(\mathcal{U}_{k}^{2} \mathbb{I}_{\mathcal{U}_{k}}\right) \leq \frac{c_{1} \log ^{2}\left(u_{r}-k+1\right)+c_{2}}{\left(u_{r}-k+1\right) \epsilon^{2}}$.

Recall that $B=\mathbb{E}\left(U_{k}^{2}\left(1-\mathbb{I}_{U_{k}}\right)\right)$. We can write $B$ as $\mathbb{E}\left(L_{k}^{2}\left(1-\mathbb{I}_{L_{k}}\right)\right)$ as $L_{k}=U_{k}$. Using the exact same arguments as above for $L_{k}$ but on sequence $-\mathbf{S}$ yields that $B \leq \frac{c_{1} \log ^{2}\left(k-l_{r}+1\right)+c_{2}}{\left(k-l_{r}+1\right) \epsilon^{2}}$.

Finally, we get that $\overline{\mathbf{S}}[k]=A+B$ which is less than $\frac{c_{1} \log ^{2}\left(u_{r}-k+1\right)+c_{2}}{\left(u_{r}-k+1\right) \epsilon^{2}}+\frac{c_{1} \log ^{2}\left(k-l_{r}+1\right)+c_{2}}{\left(k-l_{r}+1\right) \epsilon^{2}}$.

To obtain a bound on the total error $(\overline{\mathbf{S}})$.

$$
\begin{aligned}
\operatorname{error}(\overline{\mathbf{S}})= & \sum_{r=1}^{d} \sum_{k \in\left[l_{r}, u_{r}\right]} \operatorname{error}(\overline{\mathbf{S}}[k]) \\
\leq & \sum_{r=1}^{d} \sum_{k \in\left[l_{r}, u_{r}\right]} \frac{c_{1} \log ^{2}\left(u_{r}-k+1\right)+c_{2}}{\left(u_{r}-k+1\right) \epsilon^{2}}+ \\
& \sum_{r=1}^{d} \sum_{k \in\left[l_{r}, u_{r}\right]} \frac{c_{1} \log ^{2}\left(k-l_{r}+1\right)+c_{2}}{\left(k-l_{r}+1\right) \epsilon^{2}} \\
\leq & \sum_{r=1}^{d} \frac{c_{1} \log ^{3}\left(u_{r}-l_{r}+1\right)+c_{2}}{\epsilon^{2}}
\end{aligned}
$$

Finally noting that $u_{r}-l_{r}+1$ is just $n_{r}$, the number of occurrences of $s_{r}$ in $s$, we get $\operatorname{error}(\overline{\mathbf{S}})=\sum_{r} \frac{c_{1} \log ^{3} n_{r}+c_{2}}{\epsilon^{2}}=$ $O\left(d \log ^{3} n / \epsilon^{2}\right)$. This completes the proof of the theorem.

## D. 3 Proof of Theorem 3

We first restate the theorem below.
Theorem 3. Given the noisy sequence $\tilde{h}=\tilde{\mathbf{H}}(I)$, the unique minimum $L_{2}$ solution, $\bar{h}$, is given by the following recurrence relation. Let $u$ be $v$ 's parent:

$$
\bar{h}[v]=\left\{\begin{array}{l}
z[v], \quad \text { if } v \text { is the root } \\
z[v]+\frac{1}{k}\left(\bar{h}[u]-\sum_{w \in \operatorname{succ}(u)} z[w]\right), \text { o.w. }
\end{array}\right.
$$

Proof. We first show that $\bar{h}[r]=z[r]$ for the root node $r$. By definition of a minimum $L_{2}$ solution, the sequence $\bar{h}$ satisfies the following constrained optimization problem. Let $\operatorname{succ} Z[u]=\sum_{w \in \operatorname{succ}(u)} z[w]$.

$$
\begin{array}{r}
\operatorname{minimize} \\
\sum_{v}(\bar{h}[v]-\tilde{h}[v])^{2} \\
\text { subject to } \forall v, \sum_{u \in \operatorname{succ}(v)} \bar{h}[u]=\bar{h}[v]
\end{array}
$$

Denote leaves $(v)$ to be the set of leaf nodes in the subtree rooted at $v$. The above optimization problem can be rewritten as the following unconstrained minimization problem.

$$
\operatorname{minimize} \sum_{v}\left(\left(\sum_{l \in \text { leaves }(v)} \bar{h}[l]\right)-\tilde{h}[v]\right)^{2}
$$

For finding the minimum, we take derivative w.r.t $\bar{h}[l]$ for each $l$ and equate it to 0 . We thus get the following set of equations for the minimum solution.

$$
\forall l, \sum_{v: l \in l \text { leaves }(v)} 2\left(\left(\sum_{l^{\prime} \in \operatorname{leaves}(v)} \bar{h}\left[l^{\prime}\right]\right)-\tilde{h}[v]\right)=0
$$

Since $\sum_{l^{\prime} \in \operatorname{leaves}(v)} \bar{h}\left[l^{\prime}\right]=\bar{h}[v]$, the above set of equations can be rewritten as: $\forall l, \sum_{v: l \in \text { leaves }(v)} \bar{h}[v]=\sum_{v: l \in \text { leaves }(v)} \tilde{h}[v]$

For a leaf node $l$, we can think of the above equation for $l$ as corresponding to a path from $l$ to the root $r$ of the tree. The equation states that sum of the sequences $\bar{h}$ and $\tilde{h}$ over the nodes along the path are the same. We can sum all the equations to obtain the following equation.

$$
\sum_{v} \sum_{l \in \text { leaves }(v)} \bar{h}[v]=\sum_{v} \sum_{l \in \text { leaves }(v)} \tilde{h}[v]
$$

Denote level $(i)$ as the set of nodes at height $i$ of the tree. Thus root belongs to level $(\ell-1)$ and leaves in level $(0)$. Abbreviating LHS (RHS) for the left (right) hand side of the above equation, we observe the following.

$$
\begin{aligned}
L H S & =\sum_{v} \sum_{l \in \operatorname{leaves}(v)} \bar{h}[v] \\
& =\sum_{i=0}^{\ell-1} \sum_{v \in \operatorname{level}(i)} \sum_{l \in \operatorname{leaves}(v)} \bar{h}[v] \\
& =\sum_{i=0}^{\ell-1} \sum_{v \in \operatorname{level}(i)} k^{i} \bar{h}[v]=\sum_{i=0}^{\ell-1} k^{i} \sum_{v \in \operatorname{level}(i)} \bar{h}[v] \\
& =\sum_{i=0}^{\ell-1} k^{i} \bar{h}[r]=\frac{k^{\ell}-1}{k-1} \bar{h}[r]
\end{aligned}
$$

Here we use the fact that $\sum_{v \in \operatorname{level}(i)} \bar{h}[v]=\bar{h}[r]$ for any level $i$. This is because $\bar{h}$ satisfies the constraints of the tree. In a similar way, we also simplify the RHS.

$$
\begin{aligned}
R H S & =\sum_{v} \sum_{l \in \operatorname{leaves}(v)} \tilde{h}[v] \\
& =\sum_{i=0}^{\ell-1} \sum_{v \in \operatorname{level}(i)} \sum_{l \in \operatorname{leaves}(v)} \tilde{h}[v] \\
& =\sum_{i=0}^{\ell-1} \sum_{v \in \operatorname{level}(i)} k^{i} \tilde{h}[v]=\sum_{i=0}^{\ell-1} k^{i} \sum_{v \in \operatorname{level}(i)} \tilde{h}[v]
\end{aligned}
$$

Note that we cannot simplify the RHS further as $\tilde{h}[v]$ may not satisfy the constraints of the tree. Finally equating $L H S$ and $R H S$ we get the following equation.

$$
\bar{h}[r]=\frac{k-1}{k^{\ell}-1} \sum_{i=0}^{\ell-1} k^{i} \sum_{v \in \operatorname{level}(i)} \tilde{h}[v]
$$

Further, it is easy to expand $z[r]$ and check that

$$
z[r]=\frac{k-1}{k^{\ell}-1} \sum_{i=0}^{\ell-1} k^{i} \sum_{v \in \operatorname{level}(i)} \tilde{h}[v]
$$

Thus we get $\bar{h}[r]=z[r]$. For nodes $v$ other than the $r$, assume that we have computed $\bar{h}[u]$ for $u=\operatorname{pred}(v)$. Denote $H=\bar{h}[u]$. Once $H$ is fixed, we can argue that the value of $\bar{h}[v]$ will be independent of the values of $\tilde{h}[w]$ for any $w$ not in the subtree of $u$.

For nodes $w \in \operatorname{subtree}(u)$ the $L_{2}$ minimization problem is equivalent to the following one.

$$
\begin{aligned}
& \operatorname{minimize} \sum_{w \in \operatorname{subtree}(u)}(\bar{h}[w]-\tilde{h}[w])^{2} \\
& \text { subject to } \forall w \in \operatorname{subtree}(u), \sum_{w^{\prime} \in \operatorname{succ}(w)} \bar{h}\left[w^{\prime}\right]=\bar{h}[w] \\
& \text { and } \sum_{v \in \operatorname{succ}(u)} \bar{h}[v]=H
\end{aligned}
$$

Again using nodes $l \in$ leaves $(u)$, we convert this minimization into the following one.

$$
\begin{aligned}
& \operatorname{minimize} \sum_{w \in \operatorname{subtree}(U)}\left(\left(\sum_{l \in \text { leaves }(w)} \bar{h}[l]\right)-\tilde{h}[w]\right)^{2} \\
& \text { subject to } \sum_{l \in \operatorname{leaves}(u)} \bar{h}[u]=H
\end{aligned}
$$

We can now use the method of Lagrange multipliers to find the solution of the above constrained minimization problem. Using $\lambda$ as the Lagrange parameter for the constraint $\sum_{l \in \text { leaves }(u)} \bar{h}[u]=H$, we get the following sets of equations.

$$
\forall l \in \operatorname{leaves}(u), \sum_{w: l \in \operatorname{leaves}(w)} 2(\bar{h}[w]-\tilde{h}[w])=-\lambda
$$

Adding the equations for all $l \in \operatorname{leaves}(u)$ and solving for $\lambda$ we get $\lambda=-\frac{H-\operatorname{succ} Z[u]}{n(u)-1}$. Here $n(u)$ is the number of nodes in subtree( $u$ ). Finally adding the above equations for only leaf nodes $l \in \operatorname{leaves}(v)$, we get

$$
\begin{aligned}
\bar{h}[v] & =z[v]-(n(v)-1) \cdot \lambda \\
& \left.=z[v]+\frac{n(v)-1}{n(u)-1}(H-\operatorname{succ} Z[u]]\right) \\
& =z[v]+\frac{1}{k}(\bar{h}[u]-\operatorname{succ} Z[u])
\end{aligned}
$$

This completes the proof.

## D. 4 Proof of Theorem 4

First, the theorem is restated.
Theorem 4. (i) $\overline{\mathbf{H}}$ is a linear unbiased estimator, (ii) $\operatorname{error}\left(\overline{\mathbf{H}}_{q}\right) \leq \operatorname{error}\left(E_{q}\right)$ for all $q$ and for all linear unbiased estimators $E$, (iii) error $\left(\overline{\mathbf{H}}_{q}\right)=O\left(\ell^{3} / \epsilon^{2}\right)$ for all $q$, and (iv) there exists a query $q$ s.t. $\operatorname{error}\left(\overline{\mathbf{H}}_{q}\right) \leq \frac{3}{2(\ell-1)(k-1)-k} \operatorname{error}\left(\tilde{\mathbf{H}}_{q}\right)$.

Proof. For (i), the linearity of $\overline{\mathbf{H}}$ is obvious from the definition of $z$ and $\bar{h}$. To show $\overline{\mathbf{H}}$ is unbiased, we first show that $z$ is unbiased, i.e. $\mathbb{E}(z[v])=h[v]$. We use induction: the base case is if $v$ is a leaf node in which case $\mathbb{E}(z[v])=$ $\mathbb{E}(\tilde{h}[v])=h[v]$. If $v$ is not a leaf node, assume that we have shown $z$ is unbiased for all nodes $u \in \operatorname{succ}(v)$. Thus

$$
\mathbb{E}(\operatorname{succ} Z[v])=\sum_{u \in \operatorname{succ}(v)} \mathbb{E}(z[u])=\sum_{u \in \operatorname{succ}(v)} h[u]=h[v]
$$

Thus succ $Z[v]$ is an unbiased estimator for $h[v]$. Since $z[v]$ is a linear combination of $\tilde{h}[v]$ and $\operatorname{succ} Z[v]$, which are both unbiased estimators, $z[v]$ is also unbiased. This completes the induction step proving that $z$ is unbiased for all nodes. Finally, we note that $\bar{h}[v]$ is a linear combination of $\tilde{h}[v], z[v]$, and succ $Z[v]$, all of which are unbiased estimators. Thus $\bar{h}[v]$ is also unbiased proving (i).

For (ii), we shall use the Gauss-Markov theorem [21]. We shall treat the sequence $\tilde{h}$ as the set of observed variables, and $l$, the sequence of original leaf counts, as the set of unobserved variables. It is easy to see that for all nodes $v$

$$
\tilde{h}[v]=\sum_{u \in \operatorname{leaves}(v)} l[u]+\operatorname{noise}(v)
$$

Here noise $(v)$ is the Laplacian random variable, which is independent for different nodes $v$, but has the same variance
for all nodes. Hence $\tilde{h}$ satisfies the hypothesis of GaussMarkov theorem. (i) shows that $\bar{h}$ is a linear unbiased estimator. Further, $\bar{h}$ has been obtained by minimizing the $L_{2}$ distance with $\tilde{h}[v]$. Hence, $\bar{h}$ is the Ordinary Least Squares (OLS) estimator, which by the Gauss-Markov theorem has the least error. Since it is the OLS estimator, it minimizes the error for estimating any linear combination of the original counts, which includes in particular the given range query $q$.

For (iii), we note that any query $q$ can be answered by summing at most $k \ell$ nodes in the tree. Since for any node $v, \operatorname{error}(\overline{\mathbf{H}}[v]) \leq \operatorname{error}(\tilde{\mathbf{H}}[v])=2 \ell^{2} / \epsilon^{2}$, we get

$$
\operatorname{error}(\overline{\mathbf{H}}[q]) \leq k \ell\left(2 \ell^{2} / \epsilon^{2}\right)=O\left(\ell^{3} / \epsilon^{2}\right)
$$

For (iv), denote $l_{1}$ and $l_{2}$ to be the leftmost and rightmost leaf nodes in the tree. Denote $r$ to be the root. We consider the query $q$ that asks for the sum of all leaf nodes except for $l_{1}$ and $l_{2}$. Then from (i) $\operatorname{error}(\overline{\mathbf{H}}(q))$ is less than the error of the estimate $\tilde{h}[r]-\tilde{h}\left[l_{1}\right]-\tilde{h}\left[l_{2}\right]$, which is $6 \ell^{2} / \epsilon^{2}$. But, on the other hand, $\tilde{\mathbf{H}}$ will require summing $2(k-1)(\ell-1)-k$ noisy counts in total-2(k-1) at each level of the tree, except for the root and the level just below the root, only $k-2$ counts are summed. Thus $\operatorname{error}\left(\tilde{\mathbf{H}}_{q}\right)=2(2(k-1)(\ell-1)-k) \ell^{2} / \epsilon^{2}$. Thus

$$
\operatorname{error}\left(\overline{\mathbf{H}}_{q}\right) \leq \frac{3 \operatorname{error}\left(\tilde{\mathbf{H}}_{q}\right)}{2(\ell-1)(k-1)-k}
$$

This completes the proof.

## E. COMPARISON WITH BLUM ET AL.

We compare a binary $\tilde{\mathbf{H}}_{q}$ against the binary search equidepth histogram of Blum et al. [4] in terms of $(\epsilon, \delta)$-usefulness as defined by Blum et al. Since $\epsilon$ is used in the usefulness definition, we use $\alpha$ as the parameter for $\alpha$-differential privacy.

Let $N$ be the number of records in the database. An algorithm is $(\epsilon, \delta)$-useful for a class of queries if with probability at least $1-\delta$, for every query in the class, the absolute error is at most $\epsilon N$.

For any range query $q$, the absolute error of $\tilde{\mathbf{H}}_{q}$ is $\mid \tilde{\mathbf{H}}_{q}(I)-$ $\mathbf{H}_{q}(I)\left|=|Y|\right.$ where $Y=\sum_{i=1}^{c}{\underset{\tilde{\mathbf{H}}}{i}}$, each $\gamma_{i} \sim \operatorname{Lap}(\ell / \alpha)$, and $c$ is the number of subtrees in $\tilde{\mathbf{H}}_{q}$, which is at most $2 \ell$. We use Corollary 1 from [5] to bound the error of a sum of Laplace random variables. With $\nu=\sqrt{c \ell^{2} / \alpha^{2}} \sqrt{2 \ln \frac{2}{\delta^{\prime}}}$, we obtain

$$
\operatorname{Pr}\left[|Y| \leq \frac{16 \ell^{\frac{3}{2}} \ln \frac{2}{\delta^{\prime}}}{\alpha}\right] \geq 1-\delta^{\prime}
$$

The above is for a single range query. To bound the error for all $\binom{n}{2}$ range queries, we use a union bound. Set $\delta^{\prime}=\frac{\delta}{n^{2}}$. Then $\tilde{\mathbf{H}}$ is $(\epsilon, \delta)$-useful provided that $\epsilon \geq\left(16 \ell^{\frac{3}{2}} \ln \frac{2 n^{2}}{\delta}\right) / \alpha$. As in Blum et al., we can also fix $\epsilon$ and bound the size of the database. $\tilde{\mathbf{H}}$ is $(\epsilon, \delta)$-useful when

$$
N \geq \frac{16 \ell^{\frac{3}{2}} \ln \frac{2 n^{2}}{\delta}}{\alpha \epsilon}=O\left(\frac{\log ^{\frac{3}{2}} n\left(\log n+\log \frac{1}{\delta}\right)}{\alpha \epsilon}\right)
$$

In comparison, the technique of Blum et al. is $(\epsilon, \delta)$-useful for range queries when

$$
N \geq O\left(\frac{\log n\left(\log \log n+\log \frac{1}{\epsilon \delta}\right)}{\alpha \epsilon^{3}}\right)
$$

Both techniques scale at most poly-logarithmically with the size of the domain. However, the $\tilde{\mathbf{H}}$ scales better with $\epsilon$, achieving the same utility guarantee with a database that is smaller by a factor of $O\left(1 / \epsilon^{2}\right)$.

The above comparison reveals a distinction between the techniques: for $\tilde{\mathbf{H}}_{q}$ the bound on absolute error is independent of database size, i.e., it only depends on $\epsilon, \alpha$, and the size of range. However, for the Blum et al. approach, the absolute error increases with the size of the database at a rate of $O\left(N^{2 / 3}\right)$.


[^0]:    ${ }^{1}$ This is not a critical assumption and, in fact, the number

