
Differentially Private Continual Release of Histograms and Related Queries

Monika Henzinger
Institute of Science and Technology
Austria (ISTA)
Klosterneuburg, Austria

A. R. Sricharan
University of Vienna,
UniVie Doctoral School
Computer Science DoCS
Vienna, Austria

Teresa Anna Steiner
University of Southern Denmark
Odense, Denmark

Abstract

We study privately releasing column sums of a d -dimensional table with entries from a universe χ undergoing T row updates, called histogram under continual release. Our mechanisms give better additive ℓ_∞ -error than existing mechanisms for a large class of queries and input streams. Our first contribution is an output-sensitive mechanism in the insertions-only model ($\chi = \{0, 1\}$) for maintaining (i) the histogram or (ii) queries that do not require maintaining the entire histogram, such as the maximum or minimum column sum, the median, or any quantiles. The mechanism has an additive error of $O(d \log^2(dq^*) + \log T)$ whp, where q^* is the maximum output value over all time steps on this dataset. The mechanism does not require q^* as input. This breaks the $\Omega(d \log T)$ bound of prior work when $q^* \ll T$. Our second contribution is a mechanism for the turnstile model that admits negative entry updates ($\chi = \{-1, 0, 1\}$). This mechanism has an additive error of $O(d \log^2(dK) + \log T)$ whp, where K is the number of times two consecutive data rows differ, and the mechanism does not require K as input. This is useful when monitoring inputs that only vary under unusual circumstances. For $d = 1$ this gives the first private mechanism with error $O(\log^2 K + \log T)$ for continual counting in the turnstile model, improving on the $O(\log^2 n + \log T)$ error bound by [Dwork et al. \(2015\)](#), where n is the number of ones in the

stream, as well as allowing negative entries, while [Dwork et al. \(2015\)](#) can only handle nonnegative entries ($\chi = \{0, 1\}$).

1 INTRODUCTION

Maintaining continual sums of a stream of numbers is an integral subroutine in various applications, including iterative first-order methods in machine learning and online convex optimization, which require maintaining the sum of all the past gradients to make future decisions. Private mechanisms for these problems thus rely on privately maintaining continual sums of gradients ([Kairouz et al., 2021](#); [Zhang et al., 2022](#); [Denisov et al., 2022](#); [Choquette-Choo et al., 2023b](#); [Asi et al., 2023](#); [Choquette-Choo et al., 2023a](#); [Xu et al., 2023](#)), or more generally of a stream of numbers ([Dwork et al., 2010](#); [Chan et al., 2011](#); [Fichtenberger et al., 2023](#); [Henzinger et al., 2023](#); [Andersson and Pagh, 2023](#); [Henzinger and Upadhyay, 2025](#)), with minimal error. The privacy model used is the well-studied model of *differential privacy under continual observation* ([Dwork et al., 2006b, 2010](#)), which requires that the output distributions of a private mechanism are close on neighboring streams. When the elements of the stream are nonnegative, this is the *insertions-only* model, while the *turnstile model* allows negative numbers as well.

Since the initial work of [Dwork et al. \(2010\)](#) and [Chan et al. \(2011\)](#), achieving an asymptotic improvement in the additive error of private continual counting or proving that the current bound is optimal has become a major open problem in the field. Recent work has concentrated on non-asymptotic improvements: (1) [Fichtenberger et al. \(2023\)](#); [Andersson et al. \(2024\)](#); [Henzinger and Upadhyay \(2025\)](#) improved the constant term in front of the larger asymptotic terms in the additive error. (2) [Dwork et al. \(2015\)](#) improved the error to optimal on sparse streams in the

insertions-only setting with an error bound that is parameterized in the maximum output value. Currently, all improvements to the additive error are interesting since the huge error terms in existing mechanisms lead to large choices of privacy parameters in practice to ensure accurate outputs, which in turn leads to reduced privacy guarantees for these implementations. We present two improvements to a high-dimensional generalization of this problem in the spirit of the latter parameterized result of [Dwork et al. \(2015\)](#).

At a high level, [Dwork et al. \(2015\)](#) obtain their parameterized improvements by partitioning the stream of inputs into *intervals*, then batching all the inputs within an interval as a single input to a black-box continual counting mechanism. They show that the number of intervals created is proportional to the sparsity of the input stream. This reduces the error dependence on the length of the input stream to just its sparsity.

Extending continual sum of a stream of scalars to maintaining continual coordinate sums of a stream of d -dimensional vectors is a natural generalization when one needs to keep track of high-dimensional information, as is the case in the above machine learning applications. We study this setting, called the HISTOGRAM problem, which requires maintaining for every $i \in [d]$, the sum of the i -th coordinate of all (row) vectors seen so far, called the i -th *column sum*. We also show how to more accurately answer *histogram queries* that do not necessarily require maintaining the entire histogram, for example, the minimum or median column sum. Our error metric is the maximum error between the mechanism’s output and the true output, where the maximum is over all coordinates, all time steps, and all queries, called the ℓ_∞ -error.

Definition 1.1 (Continual Histogram). Let $d \in \mathbb{N}_{>0}$. The input is a stream of T row vectors x^1, \dots, x^T with $x^t \in \chi^d$ for all $t \in [T]$ where χ is the data universe and T is not given as input. In the *insertions-only* setting, $\chi = \{0, 1\}$. In the *turnstile model*, $\chi = \{-1, 0, 1\}$. The output of HISTOGRAM at each time step t is (an additive approximation to) the column sums of all inputs seen so far, i.e., $h(t) = (\sum_{\ell=1}^t x_i^\ell)_{i \in [d]}$. When $d = 1$, this is the *continual counting* problem.

An important definition in differential privacy is that of *neighboring streams*, which specifies the level of privacy guaranteed by the mechanism. We use the standard event-level neighboring definition which allows two neighboring streams to differ in all the coordinates of exactly one of the d -dimensional input vectors.

Motivating Questions We discuss ϵ -differential privacy below, and drop ϵ^{-1} and $\log(1/\beta)$ factors (β is the failure probability) in the discussion for simplicity. For private continual counting, the best known lower

bound on the additive error is $\Omega(\log T)$ ([Dwork et al., 2010](#)) against an upper bound of $O(\log^2 T)$ ([Dwork et al., 2010](#); [Chan et al., 2011](#)). In the parameterized setting, [Dwork et al. \(2015\)](#) presented a mechanism with $O(\log^2 n + \log T)$ error for insertions-only streams, where n is the largest output of the mechanism on the entire stream. Since this mechanism asymptotically improves on all other mechanisms for streams with continual count $T^{o(1)}$ and is optimal when the continual count is $O(2^{\sqrt{\log T}})$, we ask the following generalization to d -dimensional streams.

Can we exploit sparsity of outputs for private continual HISTOGRAM and histogram queries to obtain a smaller additive error?

Note that the requirement for histogram queries is stronger than that of HISTOGRAM, since the histogram query of, say, the minimum column sum has a much smaller output than that of the entire HISTOGRAM.

The best known lower bound here is $\Omega(d + \log T)$ ([Jain et al., 2023](#); [Dwork et al., 2010](#)), and we show that existing mechanisms for HISTOGRAM have an additive error of $\Omega(d \log T)$. Thus, achieving $o(d \log T)$ error necessitates new approaches. Our first mechanism improves on existing mechanisms for streams and queries with maximum output value $T^{o(1)}$, and breaks the $\Omega(d \log T)$ barrier on streams with maximum output value $o(2^{\sqrt{\log T}})$, which answers the highlighted question in the affirmative. The mechanism works by partitioning the input stream as in [Dwork et al. \(2015\)](#), while taking the interaction between all d columns into account. This extension to the multi-dimensional setting involves overcoming multiple technical difficulties which we detail later.

Another shortcoming of the mechanism of [Dwork et al. \(2015\)](#) is that it does not allow negative entries. It implicitly assumes that the continual count is increasing *monotonically*, which might not be true for turnstile streams. This motivates our second main question.

Can we obtain parameterized improvements (similar to [Dwork et al. \(2015\)](#)) for turnstile streams?

Improvements in the turnstile model are important since continual histogram mechanisms in this model are used as subroutines in more advanced private continual release mechanisms, such as in the above machine learning applications, fully dynamic graph mechanisms ([Fichtenberger et al., 2021, 2023](#)), for k -means clustering in Euclidean spaces ([la Tour et al., 2024](#)), and also to count the *difference sequence* ([Fichtenberger et al., 2021](#)) of an insertions-only stream.

We answer our second question in the affirmative, by parameterizing on the *number of fluctuations*, K , which is the number of times a row vector in the input is different from its immediately preceding row vector. We show that the improvements achieved by our first mechanism also carry over to the second, giving a mechanism for HISTOGRAM on turnstile streams that breaks the $\Omega(d \log T)$ barrier on streams with small K . At a high level, while existing mechanisms use the sparse vector technique (SVT) to update outputs when the function value changes a lot, our mechanism instead uses SVT to check when the *slope* of the function changes a lot, and we show that one can still obtain improved error bounds under this condition.

Prior Work We discuss current approaches to HISTOGRAM and related queries for ϵ -differential privacy, and present the (ϵ, δ) -dp bounds in the appendix. The bounds discussed are also presented in Tables 1 and 2.

On the lower bounds side, Jain et al. (2023) study HISTOGRAM as well as maintaining the maximum column sum (MAXSUM) and its coordinate (SUMSELECT) privately and show that the additive error¹ must be $\Omega(d)$. Combined with the counting lower bound of $\Omega(\log T)$ of Dwork et al. (2010), this gives an $\Omega(d + \log T)$ bound. Cohen et al. (2024) give a lower bound of $\Omega(\min\{n, \log T\})$ for continual counting. On the upper bound side, Jain et al. (2023) give an $O(d \log d \log^3 T)$ bound on the additive error². In the turnstile model, composing d binary tree counting mechanisms of Chan et al. (2011), and suitably scaling the privacy parameter and failure probability of each mechanism, gives an error of $O(d \log^2(dT))$. For insertion-only streams, combining the mechanism of Dwork et al. (2015) with observations by Qiu and Yi (2022) and Chan et al. (2011) gives an error bound of $O(d \log^2(dn_{\max}) + d \log T)$, where n_{\max} is the maximum column sum.

As seen from the bounds above, all existing histogram mechanisms have an $O(d \log T)$ term in their upper bounds. We show that this dependency is inherent, by proving that any histogram mechanism obtained as a composition of d counting mechanisms has $\Omega(d \log T)$ error. This implies that any mechanism that beats this barrier must necessarily take into account the interaction between the different columns.

2 OUR RESULTS

To answer the questions raised above, we present two new mechanisms. Both mechanisms achieve a param-

¹We focus only on their bounds that are subpolynomial in T since we consider the setting where $T \gg d$.

²Footnote 1 applies as well.

eterized error bound which *does not have an additive $d \log T$ term*, but is instead of the form $O(d \log^2 \rho + \log T)$, with ρ , one of the above-mentioned parameters, potentially much smaller than T .

Mechanism 1: Our first mechanism gives new parameterized upper bounds on the additive error in the insertions-only model for computing a large class of histogram-based queries that include HISTOGRAM, MAXSUM and MINSUM³, SUMSELECT, as well as QUANTILE _{p} ⁴ and TOPK⁵, where the parameter depends on the queries answered. More specifically, the additive error of the mechanism is $O(d \log^2(dq^*) + \log T)$, where q^* is the *maximum query value for the given input data at any time step*. Thus if $q^* = o(2^{\sqrt{\log T}})$, our result breaks the $\Omega(d \log T)$ bound and is better than what can be achieved by the previous mechanisms, even when the maximum column sum $n_{\max} = \Omega(T)$. Our mechanism does not need to be given q^* or T at initialization. The lower bound of $\Omega(d + \log T)$ mentioned earlier implies that the dependency on d or $\log T$ cannot be removed.

More generally, we define a class of real-valued queries, called *monotone histogram queries*, that subsumes the queries mentioned above. These queries are *monotonically increasing* when a new (nonnegative) row is added, and have low *sensitivity*. Informally, a query's *sensitivity* is the maximum difference between its output values on two neighboring streams (see Def. 3.3).

Definition 2.1 (Monotone histogram query). A function $q : \{0, 1\}^d \rightarrow \mathbb{R}_{\geq 0}$ is a *monotone histogram query* if it is monotonically increasing; has sensitivity ≤ 1 ; and depends only on the input column sums⁶.

Each individual coordinate sum satisfies this definition, as do all the queries described above. HISTOGRAM can be obtained with d monotone histogram queries, namely, each coordinate sum.

Theorem 1. Let $x = x^1, \dots, x^T$ be an insertions-only stream, and q_1, \dots, q_m be m monotone histogram queries. Mechanism 1 is ϵ -differentially private, and answers $(q_1(x^1, \dots, x^t), \dots, q_m(x^1, \dots, x^t))$ at all time steps t with a bound on the ℓ_∞ -error of

$$O((d \log^2(dm q^*/\beta) + m \log(m q^*/\beta) + \log T) \epsilon^{-1})$$

that holds with probability $\geq 1 - \beta$ simultaneously over all time steps, where $q^* = \max_{k \in [m]} q_k(x)$. Neither T nor q^* need to be given as input to the mechanism.

Answering a single monotone histogram query using the best existing ϵ -dp mechanisms requires computing

³MINSUM: Return the minimum column sum.

⁴QUANTILE _{p} : Return the p -th quantile column sum.

⁵TOPK: Return the k -th largest column sums.

⁶symmetrically, this also works for monotonically decreasing functions in the deletions-only setting.

Table 1: ℓ_∞ -error of ϵ -dp mechanisms for continual histogram queries with constant ϵ , constant failure probability in the insertions-only model, where n_{\max} is the maximum column sum, q^* is the maximum query output, and K is the number of fluctuations in the input stream. K and q^* do not need to be given to the mechanism. The binary tree result is by Dwork et al. (2010) and Chan et al. (2011), and the partitioning follows from Dwork et al. (2015); Chan et al. (2011) and Qiu and Yi (2022).

Mechanism	HISTOGRAM	m histogram queries
Jain et al. (2023)	$O(d \log d \log^3 T)$	$O(d \log d \log^3 T)$
Binary Tree	$O(d \log^2(dT))$	$O(d \log^2(dT))$
Partitioning	$O(d \log^2(dn_{\max}) + d \log T)$	$O(d \log^2(dn_{\max}) + d \log T)$
Theorem 1	$O(d \log^2(dn_{\max}) + \log T)$	$O(d \log^2(dq^*) + m \log(mq^*) + \log T)$
Theorem 2	$O(d \log^2(dK) + \log T)$	$O(d \log^2(dK) + \log T)$

Table 2: ℓ_∞ -error for ϵ -dp mechanisms with constant ϵ and constant failure probability in the turnstile model, where K is the number of fluctuations. The binary tree result is by Dwork et al. (2010) and Chan et al. (2011).

Mechanism	HISTOGRAM
Jain et al. (2023)	$O(d \log d \log^3 T)$
Binary Tree	$O(d \log^2(dT))$
Theorem 2	$O(d \log^2(dK) + \log T)$

a full histogram, which gives error $O(d \log^2(dn_{\max}) + d \log T)$. Theorem 1 improves on these bounds in three significant ways: (1) The $d \log T$ term is replaced by a $\log T$ term, giving the first mechanism that achieves a $o(d \log T)$ bound on sparse outputs. (2) The polylogarithmic dependency of $d \log^2(dn_{\max})$ on the maximum column sum n_{\max} is reduced to a polylogarithmic dependency of $d \log^2(dq^*)$ on the maximum query value q^* . For monotone histogram queries, we have $q^* \leq n_{\max}$, and q^* could be much smaller than n_{\max} , when computing, say, the minimum column sum or the median column sum. This is true, for example, on streams with a power-law distribution of the column sums. (3) We can answer up to m queries without incurring an extra additive $O(m \log T)$ error, which would happen when using standard composition.

We extend these results to natural-numbered inputs, to (ϵ, δ) -dp, and to different neighboring definitions in Section 6.

Mechanism 2: Our second mechanism gives new parameterized upper bounds in the turnstile model for continual counting and HISTOGRAM. This mechanism has an additive error of $O(d \log^2(dK) + \log T)$, where K is the number of fluctuations, which is the number of time steps where $x^t \neq x^{t+1}$. It does not need to be given K or T at initialization. This is the first improvement in the turnstile model since the 2011 bound of $O(\log^2 T)$ for continual counting (and $O(d \log^2(dT) + d \log T)$ for HISTOGRAM) by Chan et al.

(2011). Recall that the parameterized result of Dwork et al. (2015) only works for insertions-only streams.

Our result generalizes the bound of Dwork et al. (2015) in two regards: (1) to possibly negative-valued inputs, and (2) to d -dimensional inputs. Our bound improves or matches the additive error of Dwork et al. (2015) on insertions-only streams, since $K \leq 2dn_{\max}$. However, K might be considerably smaller: a data stream that contains $n/2$ ones followed by $n/2$ zeros has $K = 1$, while $n_{\max} = T/2$. Moreover, real world data often show strong time correlations, leading to a small value of K . Examples include recommendation systems, where movies or products that are popular at a given time are likely to be rated consecutively by more people, and outlier monitoring processes, where many of the generated reports are identical (when nothing special has happened).

Theorem 2. Let $x = x^1, \dots, x^T$ be a turnstile stream. Let K be the total number of times $x^t \neq x^{t+1}$, for all $t < T$. Mechanism 2 is ϵ -differentially private, and outputs an estimate of the histogram at all time steps t with ℓ_∞ -error bound $O((d \log^2(dK/\beta) + \log T) \epsilon^{-1})$ that holds with probability at least $1 - \beta$ simultaneously over all time steps. Neither T nor K need to be given as input to the mechanism.

Our result also gives new insights about stronger continual counting lower bounds: (a) Input streams consisting of $O(2^{\sqrt{\log T}})$ non-zero entries will not lead to a stronger lower bound for continual counting (Dwork et al., 2015). (b) Even further, the input streams would need to change frequently, i.e., $\omega(2^{\sqrt{\log T}})$ times, unlike the streams used in the current lower bounds of (Dwork et al., 2010; Cohen et al., 2024).

Note that tight lower bound in K , i.e., a lower bound of $\Omega(\log^2 K + \log T)$ for $d = 1$ that holds for all values of K would also imply a lower bound of $\Omega(\log^2 T)$ for continual counting (since K could be as large as T), which is a major open problem in the area.

3 PRELIMINARIES

Definition 3.1. Two streams $x = x^1, \dots, x^T$ and $y = y^1, \dots, y^T$ with $x^t, y^t \in \chi^d$ for all $t \in [T]$ are (event-level) *neighboring streams* if there exists a time step $t^* \in [T]$ such that $x^t = y^t$ for all $t \neq t^*$ and $\|x^{t^*} - y^{t^*}\|_\infty \leq 1$ for insertion-only streams, and $\|x^{t^*} - y^{t^*}\|_\infty \leq 2$ for turnstile streams.

Continual Release Mechanism A mechanism A in this model receives input $x^t \in \chi$ at every time step t , and produces an output $a^t = A(x^1, \dots, x^t)$ which may only rely on x^1 to x^t . $A^T(x) = (a^1, a^2, \dots, a^T)$ is the collection of the outputs at all time steps $\leq T$.

Definition 3.2 (Differential privacy (Dwork et al., 2006b)). A randomized mechanism A on a domain χ^T is (ϵ, δ) -differentially private $((\epsilon, \delta)$ -dp) if for all $S \in \text{range}(A^T)$ and all neighboring $x, y \in \chi^T$ we have

$$\Pr[A^T(x) \in S] \leq e^\epsilon \Pr[A^T(y) \in S] + \delta.$$

If $\delta = 0$ then A is ϵ -differentially private (ϵ -dp).

Definition 3.3 (L_p -sensitivity). The L_p -sensitivity of $f : \chi^* \rightarrow \mathbb{R}^k$ is $\max_{x, y \text{ neighboring}} \|f(x) - f(y)\|_p$. If $k = 1$, then this, the *sensitivity* of f , is equal for all p .

We use the Laplace distribution to ensure privacy.

Fact 1 (Dwork et al., 2006b). Let $f : \chi^* \rightarrow \mathbb{R}^k$ be any function with L_1 -sensitivity Δ_1 . Let $Y_i \sim \text{Lap}(\Delta_1/\epsilon)$ for $i \in [k]$. The mechanism $A(x) = f(x) + (Y_1, \dots, Y_k)$ satisfies ϵ -dp.

Fact 2. If $Y \sim \text{Lap}(b)$, then $P(|Y| \geq t \cdot b) = \exp(-t)$.

Fact 3 (Simple Composition (Dwork et al., 2006a)). Let $A_1 : \chi^* \rightarrow \text{range}(A_1)$ and $A_2 : \chi^* \times \text{range}(A_1) \rightarrow \text{range}(A_2)$ be ϵ_1 and ϵ_2 -dp mechanisms resp.. Then $A_1 \circ A_2$ is $\epsilon_1 + \epsilon_2$ -dp.

Differential Privacy Against an Adaptive Adversary As a subroutine, we use a continual histogram mechanism that works in the stronger *adaptive continual release model* defined by Jain et al. (2023). In this model, the mechanism M interacts with a *randomized adversarial process* Adv that has no restrictions on its time or space complexity. It knows the mechanism M and all its inputs and outputs up to the current time step, but *not* its random coin flips. Based on this, Adv has to choose the input to M at the next time step.

Event-level neighboring inputs are modelled as follows. All time steps except one are *regular*, and the adversary is allowed to adaptively determine when the special *challenge* time step occurs. At a regular time step, Adv outputs one value x^t . At a challenge time step, Adv outputs two values $x_{(L)}^t$ and $x_{(R)}^t$ such

that $x^1, \dots, x_{(L)}^t, x^{t+1}, \dots$ and $x^1, \dots, x_{(R)}^t, x^{t+1}, \dots$ are neighboring (here, $\|x_{(L)}^t - x_{(R)}^t\|_\infty \leq 1$). At the challenge time step, an external oracle selects one of these two inputs and sends it to M . The oracle decides at the beginning of the interaction whether it sends the first or the second input to M . Importantly, this decision is not known either to Adv or M . The goal of the adversary is to determine which decision was made by the oracle, while the goal of the mechanism is to return the computed output, e.g., a histogram, such that Adv does not find out which decision was made by the oracle. A more formal description of the adaptive continual release model is given in the appendix. Denisov et al. (2022) shows that ϵ -dp under continual release implies ϵ -dp against an adaptive adversary, which is not true for (ϵ, δ) -dp in general.

Fact 4 (Prop 2.1 of Denisov et al. (2022)). Every mechanism that is ϵ -differentially private in the continual release model is also ϵ -dp in the adaptive continual release model.

Continual Counting and Histogram The additive error bound of the ϵ -dp continual counting mechanism by Chan et al. (2011) is $O(\log^2(T/\beta)\epsilon^{-1})$ with probability $1 - \beta$ for *all time steps* $t \in [T]$ *simultaneously*, where T is not given as input to the mechanism. The inputs are allowed to be integers, and neighboring is as in Definition 3.1. A simple mechanism for continual histogram is using d binary counting mechanisms, one per column. With Facts 3 and 4, this yields:

Fact 5. There is an ϵ -dp mechanism for continual histogram in the adaptive model whose ℓ_∞ -error is bounded by $O(d \log^2(dT/\beta)\epsilon^{-1})$ with probability $1 - \beta$.

This gives a blackbox reduction from ϵ -dp continual histogram to ϵ -dp continual counting. We show in Appendix D that any continual histogram mechanism constructed this way must have an error of $\Omega(d \log T)$:

Lemma 3.1. *Let A be any (ϵ/d) -dp continual counting mechanism. Then the histogram mechanism H , defined by running A independently for each coordinate, must have an error of $\Omega(d \log(T)\epsilon^{-1})$ at some time step $t \leq T$ with constant probability.*

4 HISTOGRAM QUERIES PARAMETERIZED IN MAXIMUM QUERY OUTPUT

Mechanism 1 is designed to answer m monotone histogram queries with output-sensitive error on insertion-only streams. It consists of two main parts, a *partitioning* mechanism and a black-box *histogram* mechanism H . The goal of the partitioning mecha-

nism is to sparsify the input stream provided to the histogram mechanism H by partitioning the input stream into *intervals*. It batches consecutive inputs to Mechanism 1 together into an interval, and combines these inputs into a single input to H .

In more detail, the algorithm keeps parameters c_i and s_i for each coordinate, which keep the current estimate of the column sum within the current interval, and the total column sum, respectively. On an input x^t , it updates these values (line 14). It then tests if any of the queries on s crosses a threshold (line 23). If not, it returns the output from the previous round. If it does, we insert the c_i values into the blackbox histogram algorithm H (line 25). It then updates the thresholds, the parameters s_i , and the output (lines 34-37). Note that we keep a separate threshold for each query, and they are updated differently depending on whether or not the query answer was close to the threshold in this round (lines 30-31).

The parameters and thresholds are chosen to minimize the additive error: The longer the intervals, the smaller the error from the histogram mechanism (since it has fewer insertions), and the larger the error *within* an interval (since the same output is used for all time steps within an interval). The parameters of the mechanism (Thresh_k^t , D_j^t , and C_j^t for example) are chosen with the goal of balancing these two kinds of error.

We want to point out two main differences in our approach compared to previous work, which are due to two difficulties: first, we compute non-linear queries on high-dimensional input data, and second, we want to break the $\Omega(d \log T)$ barrier. We explain first the issues that arise and then how we overcome them.

We explain the first difficulty for the query MINSUM, but it applies correspondingly to other monotone histogram queries as well. When $d = 1$ (continual counting as in Dwork et al. (2015)), the partitioning mechanism does not depend on the output of the black-box counting/histogram mechanism. This is because the continual count over a stream is equal to the sum of the continual counts of all intervals. This does not hold for non-linear queries on higher dimensional inputs because the input stream cannot be decomposed into intervals for queries like MINSUM, i.e., the MINSUM of the entire stream cannot be obtained from knowing just the MINSUM value of each interval. Instead, the partitioning algorithm requires an estimate of the current MINSUM value at every time step.

Since we do not want the computation during the current interval to depend on the private data from prior intervals, we reuse the last output of H , as it is a privatized number, in order to keep a running estimate of MINSUM. This, however, leads to the following tech-

Mechanism 1 Mechanism for answering m histogram queries parameterized in the maximum query output.

```

1: Input: Stream  $x^1, x^2, \dots \in \{0, 1\}^d$ , an adaptively
    $\epsilon$ -differentially private continual histogram mechanism  $H$ , failure probability  $\beta$ , additive error bound
    $\text{err}(t, \beta)$  that holds with probability  $\geq 1 - \beta$  for the
   output of  $H$  at time step  $t$ .
2: Output: Estimate of  $q_k(h(t))$  for all  $k \in [m]$  and
   all  $t \in \mathbb{N}$ 
3:  $\triangleright$  Initialization  $\triangleleft$ 
4: Initialize  $H$ 
5:  $\beta' = 6\beta/\pi^2$ ,  $\beta_t = \beta'/t^2$  for any  $t \in \mathbb{N}$ 
6:  $\text{Thresh}_k^1 \leftarrow 3\epsilon^{-1}(12 \ln(2/\beta') + 6 \ln(6/\beta') +$ 
    $m \ln(6m/\beta')) + 3 \cdot \text{err}(1, \beta'/6)$  for all  $k \in [m]$ 
7:  $c_i \leftarrow 0$  for all  $i \in [d]$   $\triangleright$  column sum within interval
8:  $s_i \leftarrow 0$  for all  $i \in [d]$   $\triangleright$  histogram estimate
9:  $j \leftarrow 1$   $\triangleright$  number of intervals
10:  $\tau_1 \leftarrow \text{Lap}(6/\epsilon)$ 
11:  $\text{out} \leftarrow (q_1(\mathbf{0}), q_2(\mathbf{0}), \dots, q_m(\mathbf{0}))$   $\triangleright$  current output
12:  $\triangleright$  Process the input stream  $\triangleleft$ 
13: for  $t \in \mathbb{N}$  do
14:    $c_i \leftarrow c_i + x_i^t$ ,  $s_i \leftarrow s_i + x_i^t$  for all  $i \in [d]$ 
15:    $\triangleright$  Set parameters  $\triangleleft$ 
16:    $\alpha_\mu^t \leftarrow 12\epsilon^{-1} \ln(2/\beta_t)$ 
17:    $\alpha_\tau^j \leftarrow 6\epsilon^{-1} \ln(6/\beta_j)$ 
18:    $\alpha_\gamma^j \leftarrow 3\epsilon^{-1} m \ln(6m/\beta_j)$ 
19:    $\alpha_H^j \leftarrow \text{err}(j, \beta_j/6)$ 
20:    $C_j^t \leftarrow \alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j$ ,  $D_j^t \leftarrow 3(C_j^t + \alpha_H^j)$ 
21:    $\triangleright$  Test for threshold  $\triangleleft$ 
22:    $\mu_t \leftarrow \text{Lap}(12/\epsilon)$ 
23:   if  $\exists k \in [m] : q_k(s) + \mu_t > \text{Thresh}_k^t + \tau_j$  then
24:      $\triangleright$  Close the current interval  $\triangleleft$ 
25:     insert  $(c_1, \dots, c_d)$  into  $H$ 
26:     reset  $c_i \leftarrow 0$  for all  $i \in [d]$ 
27:     for  $k \in [m]$  do
28:        $\gamma_k^j \leftarrow \text{Lap}(3m/\epsilon)$ 
29:        $\triangleright$  if  $q_k(s)$  is "close" to threshold, in-
         crease threshold  $\triangleleft$ 
30:       if  $q_k(s) + \gamma_k^j > \text{Thresh}_k^t - C_j^t$  then
31:          $\text{Thresh}_k^t \leftarrow \text{Thresh}_k^t + D_j^t$ 
32:        $j \leftarrow j + 1$ 
33:        $\triangleright$  update threshold for the new interval  $\triangleleft$ 
34:        $\text{Thresh}_k^t \leftarrow \text{Thresh}_k^t - D_{j-1}^t + D_j^t \forall k \in [m]$ 
35:        $\tau_j \leftarrow \text{Lap}(6/\epsilon)$   $\triangleright$  pick fresh noise
36:        $(s_1, \dots, s_d) \leftarrow \text{output}(H)$ 
37:        $\text{out} \leftarrow (q_1(s), \dots, q_m(s))$ 
38:     output  $\text{out}$ 
39:    $\text{Thresh}_k^{t+1} \leftarrow \text{Thresh}_k^t - D_j^t + D_j^{t+1} \forall k \in [m]$ 

```

nical challenge: The partitioning mechanism *depends on the outputs of the histogram mechanism of prior intervals*, and the input to the histogram mechanism depends on the output of the partitioning mechanism,

and, hence, on the prior output of the histogram mechanism. Furthermore, given two neighboring streams to Mechanism 1, the input to the black-box histogram mechanism that is generated by the partitioning mechanism might not necessarily be neighboring streams (consider the case where two wildly different partitions are created on two neighboring streams, leading to inputs to the histogram mechanism that are very far apart). Thus, we cannot use a simple composition theorem to show privacy for the combined mechanism.

To overcome this difficulty, we use a continuous histogram mechanism that is differentially private *even if the inputs are chosen adaptively*. We note that *concurrent* composition theorems as given in, e.g., Haney et al. (2023), cannot be used here in a black-box manner, and explain the reasons in more detail in Appendix E. Adaptive differential privacy of the continuous histogram mechanism allows us to separate the privacy loss incurred by the partitioning mechanism from that of the histogram mechanism.

The second difficulty relates to the $\Omega(d \log T)$ barrier. A naïve partitioning technique is to maintain independent thresholds for each query, and to use the sparse vector technique (SVT) separately for each query to check if the query answer is larger than its threshold. If a query answer is larger than the threshold, an interval is closed and the thresholds are increased accordingly. Since this involves interacting with private data at *each* time step for *every* query, this approach incurs the $\Omega(d \log T)$ barrier of all prior work.

To overcome this, we design a partitioning mechanism that works as follows: it checks first if *there exists* a query that crosses a certain predefined threshold value (line 23). If not, then the mechanism adds the current input to the batch and *does not close* the current interval. The output that was used for the previous time step is reused.

If there exists a query crossing the threshold, then the mechanism *closes* the current interval, sends the batched input to H , and initializes a new interval. At this point, the mechanism has privately determined that *there exists* a query that crosses the threshold. However, this information is not enough to update the thresholds, since we need to also need to privately determine the *identities* of all the queries that cross the thresholds, which we do next.

Finally, the mechanism checks each query independently and privately if its threshold needs to be updated (line 30). The parameters are chosen such that at least one query has its threshold updated at the end of each interval.

For the utility proof of Dwork et al. (2015), it was

enough to show that their choice of threshold was larger than the standard deviation of their Laplace random variables. Due to the interplay between several submechanisms, our utility proof requires a more involved analysis.

Theorem 3. *Let H be any $(\epsilon/3)$ -dp continual histogram mechanism with $\epsilon > 0$, and let q_1, \dots, q_m be any m monotone histogram queries. Mechanism 1 satisfies ϵ -dp. If H is the mechanism from Fact 5, then on input x , Mechanism 1 has additive error*

$$O((d \log^2(dm q^*/\beta) + m \log(m q^*/\beta) + \log t) \epsilon^{-1})$$

at all $t \in [T]$ simultaneously with probability $1 - \beta$, where $q^ = \max_{k \in [m]} \max_{t \in [T]} q_k(x^1, x^2, \dots, x^t)$.*

5 HISTOGRAM PARAMETERIZED IN THE NUMBER OF FLUCTUATIONS

Mechanism 2 is designed to answer HISTOGRAM in the turnstile model with a better error bound when the number of times two consecutive rows differ is small, called the *number of fluctuations*, K . The high-level structure of the mechanism is similar to that of Mechanism 1, consisting of a partitioning mechanism that interacts with a black-box histogram mechanism H . The partitioning mechanism batches inputs together into an *interval*, and sends the combined inputs in an interval as a single input to the histogram mechanism.

Earlier, the output of Mechanism 1 in an interval was set to the previous output of H , and the output remained unchanged within the interval. In general, when balancing the privacy-accuracy trade-off by limiting the number of time steps for which the private data is accessed, the classic strategy (as by Dwork et al. (2015)) is to not update the output at all between two updates to the black-box mechanism (here, the mechanism H). We go beyond this paradigm with Mechanism 2, by modifying the estimate *even within an interval*. In particular, our histogram estimate is the sum of the last output of H and a function that is linear in the length of the interval. The function is chosen such that if the input stream remains *stable*, i.e. the value does not change often, we do not need to end the current interval often (here, only $O(K)$ times).

Specifically, the output of Mechanism 2 within an interval mimics the behavior of the previous batch of updates. If the histogram was “significantly” increasing (or decreasing) for a coordinate during the previous interval, then the output of Mechanism 2 for this coordinate in the current interval adds (or subtracts) an estimate to the last output of H at each time step. This corresponds to guessing the *first-order derivative* or gradient of each coordinate of the histogram

Mechanism 2 Mechanism for HISTOGRAM parameterized in the number of fluctuations.

```

1: Input: Stream  $x^1, x^2, \dots \in \{-1, 0, 1\}^d$ , an  $\epsilon/3$ -
   differentially private continual histogram mechanism  $H$ , failure probability  $\beta$ , additive error bound
    $\text{err}(t, \beta)$  that holds with probability  $\geq 1 - \beta$  for the
   output of  $H$  at time step  $t$ .
2: Output: Estimate  $h(t)$  at all  $t \in \mathbb{N}$ 
3:  $\triangleright$  Initialization of all parameters  $\triangleleft$ 
4: Initialize  $H$ 
5:  $\beta' = 6\beta/\pi^2$ ,  $\beta_t = \beta'/t^2$  for any  $t \in \mathbb{N}$ 
6:  $j \leftarrow 1$   $\triangleright$  number of intervals
7:  $c_i \leftarrow 0$  for all  $i \in [d]$   $\triangleright$  column sum within interval
8:  $\text{mode}_i \leftarrow 0$  for all  $i \in [d]$ 
9:  $t_{\text{diff}} \leftarrow 0$   $\triangleright$  length of current interval
10:  $\tau_1 \leftarrow \text{Lap}(9/\epsilon)$ ,  $\tau_2 \leftarrow \text{Lap}(9/\epsilon)$ 
11:  $H_{\text{out}} = 0^d$   $\triangleright$  initial histogram
12:  $\triangleright$  Process the input stream  $\triangleleft$ 
13: for  $t \in \mathbb{N}$  do
14:    $c_i \leftarrow c_i + x_i^t$  for all  $i \in [d]$ 
15:    $t_{\text{diff}} \leftarrow t_{\text{diff}} + 1$ 
16:    $\alpha_t \leftarrow \frac{27}{\epsilon} \log(4/\beta_t) + \frac{3d}{\epsilon} \log(1/\beta_j)$ 
17:    $\text{Thresh}_{i,1} \leftarrow \text{mode}_i \cdot t_{\text{diff}} - 2\alpha_t$ ,  $i \in [d]$ 
18:    $\text{Thresh}_{i,2} \leftarrow \text{mode}_i \cdot t_{\text{diff}} + 2\alpha_t$ ,  $i \in [d]$ 
19:    $\mu_1^t \leftarrow \text{Lap}(18/\epsilon)$ 
20:    $\mu_2^t \leftarrow \text{Lap}(18/\epsilon)$ 
21:   if  $\min_i (c_i - \text{Thresh}_{i,1}) < \tau_1 - \mu_1^t$  then
22:      $\triangleright$  close the current interval  $\triangleleft$ 
23:     insert  $(c_1, \dots, c_d)$  into  $H$ 
24:      $H_{\text{out}} \leftarrow \text{output}(H)$   $\triangleright$  update histogram
25:     for  $i \in [d]$  do
26:        $\lambda_i = \text{Lap}(3d/\epsilon)$ 
27:        $\triangleright$  update modes  $\triangleleft$ 
28:       if  $c_i + \lambda_i < \text{Thresh}_{i,1} + \alpha_t$  then
29:          $\text{mode}_i \leftarrow \max\{\text{mode}_i - 1, -1\}$ 
30:       reset  $c_i \leftarrow 0$  for all  $i \in [d]$ 
31:        $j \leftarrow j + 1$ 
32:        $t_{\text{diff}} \leftarrow 0$ ;  $\tau_1 \leftarrow \text{Lap}(9/\epsilon)$ 
33:   else if  $\max_i (c_i - \text{Thresh}_{i,2}) > \tau_2 - \mu_2^t$  then
34:      $\triangleright$  close the current interval  $\triangleleft$ 
35:     insert  $(c_1, \dots, c_d)$  into  $H$ 
36:      $H_{\text{out}} \leftarrow \text{output}(H)$   $\triangleright$  update histogram
37:     for  $i \in [d]$  do
38:        $\lambda_i = \text{Lap}(3d/\epsilon)$ 
39:        $\triangleright$  update modes  $\triangleleft$ 
40:       if  $c_i + \lambda_i > \text{Thresh}_{i,2} - \alpha_t$  then
41:          $\text{mode}_i \leftarrow \min\{\text{mode}_i + 1, 1\}$ 
42:       reset  $c_i \leftarrow 0$  for all  $i \in [d]$ 
43:        $j \leftarrow j + 1$ 
44:        $t_{\text{diff}} \leftarrow 0$ ;  $\tau_2 \leftarrow \text{Lap}(9/\epsilon)$ 
45:   output  $H_{\text{out}} + \text{mode} \cdot t_{\text{diff}}$ 

```

within the previous interval and then using this to vary the output of the mechanism for the current interval. When the additive error of the estimate accumulated within an interval crosses a pre-specified threshold for at least one coordinate, the current interval is closed and the guess for the next interval is recomputed based on whether the prior guess overestimated (or underestimated) the true count for the current interval.

In detail, the variable c_i tracks the true count for coordinate i within the current interval, and t_{diff} is used to count the number of time steps in the interval. Crucially, we introduce the variables mode_i for $i \in [d]$, which assume values in $\{-1, 0, 1\}$ and are used to guess the slope for each coordinate. Each mode_i is initially 0. Upon an input x^t , we first update c_i and t_{diff} (lines 14-15) and the thresholds (lines 17-18). We use $\text{mode}_i \cdot t_{\text{diff}}$ as a guess for the count within the current interval for each coordinate i – that is, if $\text{mode}_i = 1$ (or 0 or -1) we guess that coordinate i consists in this interval only of 1's (or 0's or -1 's). In each round we check if for some i , the guess is too large or too small compared to c_i (lines 21 and 33). If not, we output the previous histogram output plus $\text{mode}_i \cdot t_{\text{diff}}$ for each coordinate i . If for some i , the guess is too large, we insert the c_i values into the blackbox histogram algorithm H and update its output (lines 23-24). We then reduce the modes of all coordinates i where $\text{mode}_i \cdot t_{\text{diff}}$ is too large compared to c_i (line 29). If the guess is too small for some i , we perform analogous updates (lines 33-44).

For our parameterized error bound, we use the following novel analysis approach: we subdivide the input stream into maximal contiguous substreams during which the input does not change, called *episodes*, creating at most $K + 1$ episodes in the entire stream. Within each episode, we prove that at most 9 intervals are closed with high probability (whp) as follows: the increase of c_i in one time step, called the *slope of coordinate i* , is constant within an episode and belongs to $\{-1, 0, 1\}$. Thus, within an episode, all d coordinates can be placed into three groups according to their slope. Within each *interval*, the mechanism maintains one of three different modes for each coordinate. We first show that for all coordinates with an identical slope and an identical mode, their modes are updated at the same time step whp. As intervals are only closed when the mode of some coordinate is modified, it suffices to bound the number of time steps within an episode when a mode is updated.

Whenever the mode of a coordinate changes, it holds whp that the absolute difference of its mode and its slope is reduced by 1. Further, the mode will not be updated anymore in this episode when it matches the slope of the coordinate. Thus, if the slope of a group is, say, 1, and there is a subgroup of coordinates of the

group with mode -1 at the beginning of the episode, then there will be at most two time steps where the mode of this subgroup changes, namely first to 0 and then to 1 . Counting all subgroups and including the potential interval closure when the episode ends results in up to nine intervals closed within an episode whp.

Thus, taking first-order information into account admits improved bounds parameterized in the number of fluctuations while also handling *negative* inputs, and the novel extension of SVT allows the first use of input partitioning techniques in the turnstile model.

Theorem 4. *Let H be any $(\epsilon/3)$ -dp continual histogram mechanism with $\epsilon > 0$. Mechanism 2 satisfies ϵ -dp. If H is the mechanism from Fact 5, then on input x , Mechanism 2 has additive error $O((d \log^2(dK/\beta) + \log T)\epsilon^{-1})$ at all $t \in [T]$ simultaneously with probability $1 - \beta$, where $K = \sum_{t \in [T]} \mathbb{1}(x_t \neq x_{t-1})$.*

6 EXTENSIONS

In this section, we present extensions and applications of the results shown earlier.

In the case of Mechanism 1, our bounds also extend to the setting when the entries are natural numbers ($\chi = \mathbb{N}$). We produce our bounds for (ϵ, δ) -dp in Appendix F, where the linear dependencies on d and m are replaced by \sqrt{d} and \sqrt{m} respectively, achieving a similar improvement over prior work as for ϵ -dp. In the standard definition, neighboring streams x and y may differ in the *entire* input vector at one time step, i.e., there is one time step t^* such that x^{t^*} and y^{t^*} could differ in all d coordinates. If x^{t^*} and y^{t^*} may only differ in up to $b < d$ coordinates, or, more generally, $\|x^{t^*} - y^{t^*}\|_1 \leq b$, then the linear dependency on d is replaced by the same dependency in b .

We summarize the results of Theorem 1 applied to some widely used query functions:

Corollary 6.1. *Let $x = x^1, \dots, x^T$ be an insertions-only stream as in Definition 1.1. Consider the queries HISTOGRAM, MAXSUM, SUMSELECT, TOPK, MEDIAN, MINSUM. Mechanism 1 answers the query at all time steps t , is ϵ -dp, and has an ℓ_∞ -error of $O((d \log^2(d\rho/\beta) + \log T)\epsilon^{-1})$ with probability $\geq 1 - \beta$ simultaneously over all time steps, where the parameter ρ is n_{\min} for MINSUM, n_{median} for MEDIAN, and n_{\max} for the rest, where n_{\min} , n_{median} , and n_{\max} are the minimum, median, and maximum column sum.*

As in Mechanism 1, the same technique replaces the d term with a \sqrt{d} term for (ϵ, δ) -dp for Mechanism 2 as shown in Appendix F. Similarly, if two neighboring streams may differ only in up to $b < d$ coordinates at one time step, then the linear dependency in d gets replaced by b .

7 CONCLUSION

We have presented black-box reductions from continual counting to continual histogram via partitioning mechanisms, and any improvement to continual counting immediately leads to improvements for our parameterized mechanisms as well.

In addition to providing stronger upper bounds on sparse queries and streams, our results further provide the following insights into possible approaches for stronger lower bounds for these problems. In the case of continual counting, the current lower bound of $\Omega(\log T)$ is shown using sequences that only have *two* switches. Thus, Algorithm 2 (as well as the algorithm of Dwork et al. (2015)) give an error of $O(\log T)$ on these sequences. However, the latter algorithm only works for inputs in $\{0, 1\}^T$, while ours works even for inputs from $\{-1, 0, 1\}^T$. It follows that in order to show a stronger lower bound for continual counting (if it exists) a sequence with a large number of switches has to be used, even if the input is from $\{-1, 0, 1\}^T$.


In the case of continual histograms, we show that while existing algorithms must have an $\Omega(d \log T)$ error, there is hope of removing this $d \log T$ dependency. On the hard sequences which achieve the $\Omega(d \log T)$ lower bound for existing algorithms, Algorithm 2 achieves an error of $O(d \log^2 d + \log T)$ and Algorithm 1 achieves an error of $O(d \log^2 d + d \log \log T + \log T)$.

8 FUTURE DIRECTIONS

A major open question is to close the $O(\log^2 T)$ vs $\Omega(\log T)$ gap for continual counting and ϵ -differential privacy. Additionally, for continual histogram, it would be very interesting to see if there exists an ϵ -dp algorithm with $\tilde{O}(d + \log^2 T)$ error, separating the dependence of the multiple dimensions from the dependence on $\log T$ for *all* streams - or, if a lower bound of $\Omega(d \log T)$ for *all* algorithms for continual histogram exists.

Next, many histogram settings contain streams drawn from a specific underlying probability distribution. One could imagine a histogram algorithm that *learns* from its output history, and predicts future histogram values, only updating the histogram when the error of the prediction is too large. This would extend our parameterized results to the domain of *learning-augmented private algorithms*, which could reduce the observed error by a large factor in practice. Often, observed data is highly structured and not adversarial, admitting much lower error bounds than for the adversarial case.

Acknowledgements

MH: This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (MoDynStruct, No. 101019564)  and the Austrian Science Fund (FWF) grant DOI 10.55776/Z422, grant DOI 10.55776/I5982, and grant DOI 10.55776/P33775 with additional funding from the netidee SCIENCE Stiftung, 2020–2024. TAS: This work was supported by a research grant (VIL51463) from VILLUM FONDEN.

References

- J. D. ANDERSSON AND R. PAGH. [A smooth binary mechanism for efficient private continual observation](#). In *Neural Information Processing Systems (NeurIPS)*, 2023.
- J. D. ANDERSSON, R. PAGH, AND S. TORKAMANI. [Count on your elders: Laplace vs gaussian noise](#). *arXiv*, 2024.
- H. ASI, V. FELDMAN, T. KOREN, AND K. TALWAR. [Near-optimal algorithms for private online optimization in the realizable regime](#). In *International Conference on Machine Learning (ICML)*, 2023.
- T. H. CHAN, E. SHI, AND D. SONG. [Private and continual release of statistics](#). *ACM Trans. Inf. Syst. Secur.*, 14(3):26:1–26:24, 2011.
- C. A. CHOQUETTE-CHOO, A. GANESH, R. MCKENNA, H. B. MCMAHAN, J. RUSH, A. G. THAKURTA, AND Z. XU. [\(Amplified\) banded matrix factorization: A unified approach to private training](#). In *Neural Information Processing Systems (NeurIPS)*, 2023a.
- C. A. CHOQUETTE-CHOO, H. B. MCMAHAN, J. K. RUSH, AND A. G. THAKURTA. [Multi-epoch matrix factorization mechanisms for private machine learning](#). In *International Conference on Machine Learning (ICML)*, 2023b.
- E. COHEN, X. LYU, J. NELSON, T. SARLÓS, AND U. STEMMER. [Lower bounds for differential privacy under continual observation and online threshold queries](#). In *Conference on Learning Theory (COLT)*, 2024.
- S. DENISOV, H. B. MCMAHAN, J. RUSH, A. D. SMITH, AND A. G. THAKURTA. [Improved differential privacy for SGD via optimal private linear operators on adaptive streams](#). In *Neural Information Processing Systems (NeurIPS)*, 2022.
- C. DWORK AND A. ROTH. [The algorithmic foundations of differential privacy](#). *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014.
- C. DWORK, K. KENTHAPADI, F. MCSHERRY, I. MIRONOV, AND M. NAOR. [Our data, ourselves: Privacy via distributed noise generation](#). In *Advances in Cryptology (EUROCRYPT)*, 2006a.
- C. DWORK, F. MCSHERRY, K. NISSIM, AND A. D. SMITH. [Calibrating noise to sensitivity in private data analysis](#). In *Theory of Cryptography Conference (TCC)*, 2006b.
- C. DWORK, M. NAOR, O. REINGOLD, G. N. ROTHBLUM, AND S. P. VADHAN. [On the complexity of differentially private data release: efficient algorithms and hardness results](#). In *Symposium on Theory of Computing (STOC)*, 2009.
- C. DWORK, M. NAOR, T. PITASSI, AND G. N. ROTHBLUM. [Differential privacy under continual observation](#). In *Symposium on Theory of Computing (STOC)*, 2010.
- C. DWORK, M. NAOR, O. REINGOLD, AND G. N. ROTHBLUM. [Pure differential privacy for rectangle queries via private partitions](#). In *Advances in Cryptology (ASIACRYPT)*, 2015.
- H. FICHTENBERGER, M. HENZINGER, AND L. OST. [Differentially private algorithms for graphs under continual observation](#). In *European Symposium on Algorithms (ESA)*, 2021.
- H. FICHTENBERGER, M. HENZINGER, AND J. UPADHYAY. [Constant matters: Fine-grained error bound on differentially private continual observation](#). In *International Conference on Machine Learning (ICML)*, 2023.
- P. GUERRA-BALBOA, À. MIRANDA-PASCUAL, J. PARRA-ARNAU, AND T. STRUFE. [Composition in differential privacy for general granularity notions](#). In *Computer Security Foundations Symposium (CSF)*, 2024.
- S. HANEY, M. SHOEMATE, G. TIAN, S. P. VADHAN, A. VYRROS, V. XU, AND W. ZHANG. [Concurrent composition for interactive differential privacy with adaptive privacy-loss parameters](#). In *Conference on Computer and Communications Security (CCS)*, 2023.
- M. HARDT AND K. TALWAR. [On the geometry of differential privacy](#). In *Symposium on Theory of Computing (STOC)*, 2010.
- M. HENZINGER AND J. UPADHYAY. [Improved differentially private continual observation using group algebra](#). In *Symposium on Discrete Algorithms (SODA)*, 2025.
- M. HENZINGER, J. UPADHYAY, AND S. UPADHYAY. [Almost tight error bounds on differentially private continual counting](#). In *Symposium on Discrete Algorithms (SODA)*, 2023.

- P. JAIN, S. RASKHODNIKOVA, S. SIVAKUMAR, AND A. D. SMITH. **The price of differential privacy under continual observation**. In *International Conference on Machine Learning (ICML)*, 2023.
- P. KAIROUZ, B. MCMAHAN, S. SONG, O. THAKKAR, A. THAKURTA, AND Z. XU. **Practical and private (deep) learning without sampling or shuffling**. In *International Conference on Machine Learning (ICML)*, 2021.
- M. D. LA TOUR, M. HENZINGER, AND D. SAULPIC. **Making old things new: A unified algorithm for differentially private clustering**. In *International Conference on Machine Learning (ICML)*, 2024.
- M. LYU, D. SU, AND N. LI. **Understanding the sparse vector technique for differential privacy**. *Proc. VLDB Endow.*, 10(6):637–648, 2017.
- Y. QIU AND K. YI. **Differential privacy on dynamic data**. *arXiv*, 2022.
- S. P. VADHAN AND T. WANG. **Concurrent composition of differential privacy**. In *Theory of Cryptography Conference (TCC)*, 2021.
- Z. XU, Y. ZHANG, G. ANDREW, C. A. CHOQUETTE-CHOO, P. KAIROUZ, H. B. MCMAHAN, J. ROSENSTOCK, AND Y. ZHANG. **Federated learning of gboard language models with differential privacy**. In *Meeting of the Association for Computational Linguistics (ACL): Industry Track*, 2023.
- Q. ZHANG, H. TRAN, AND A. CUTKOSKY. **Differentially private online-to-batch for smooth losses**. In *Neural Information Processing Systems (NeurIPS)*, 2022.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Not Applicable]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Not Applicable]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A PRELIMINARIES (APPENDIX)

We present the adaptive adversarial model in more detail below, and recall some useful results for our proofs.

Notation We denote the set $\{1, \dots, n\}$ by $[n]$.

Continual Release Model In the continual release model, at every time step t , we add an element $x^t \in \chi^d$ to the current data set. Note that for simplicity, we define the model as adding x^t to the data set; in the turnstile model we allow x^t to have negative entries, capturing both insertions and deletions. The entire stream of insertions is of length T , and T is not given as input to the mechanism.

Definition A.1 (Laplace Distribution). The *Laplace distribution* centered at 0 with scale b is the distribution with probability density function

$$f_{\text{Lap}(b)}(x) = \frac{1}{2b} \exp\left(\frac{-|x|}{b}\right).$$

We use $X \sim \text{Lap}(b)$ or just $\text{Lap}(b)$ to denote a random variable X distributed according to $f_{\text{Lap}(b)}(x)$.

Differential Privacy Against an Adaptive Adversary As a subroutine, we use a continual histogram mechanism that works in the stronger *adaptive continual release model* defined by Jain et al. (2023). In this model, the mechanism M interacts with a *randomized adversarial process* Adv that has no restrictions on its time or space complexity. It knows the mechanism M and all its inputs and outputs up to the current time step, but *not* its random coin flips. Based on this, Adv has to choose the input to M for the next time step.

Event-level neighboring inputs are modelled as follows. All time steps except one are *regular*, and the adversary is allowed to adaptively determine when the special *challenge* time step occurs. At a regular time step, Adv outputs one value x^t . At a challenge time step, Adv outputs two values $x_{(L)}^t$ and $x_{(R)}^t$ such that $x^1, \dots, x_{(L)}^t, x^{t+1}, \dots$ and $x^1, \dots, x_{(R)}^t, x^{t+1}, \dots$ are neighboring (here, $\|x_{(L)}^t - x_{(R)}^t\|_\infty \leq 1$). At the challenge time step, an external oracle selects one of these two inputs and sends it to M . The oracle decides before the beginning of the interaction whether it sends the first or the second input to M . Importantly, this decision is not known either to Adv or M . The goal of the adversary is to determine which decision was made by the oracle, while the goal of the mechanism is to return the computed output, e.g., a histogram, such that Adv does not find out which decision was made by the oracle.

Game 3 Privacy game $\Pi_{M, Adv}$ for the adaptive continual release model

- 1: **Input:** Stream length $T \in \mathbb{N}$, side $\in \{L, R\}$ (not known to Adv and M)
 - 2: **for** $t \in [T]$ **do**
 - 3: Adv **outputs** $\text{type}^t \in \{\text{challenge}, \text{regular}\}$, where **challenge** is only chosen for exactly one value of t
 - 4: **if** $\text{type}^t = \text{regular}$ **then**
 - 5: Adv **outputs** $x^t \in \chi$ which is sent to M
 - 6: **if** $\text{type}^t = \text{challenge}$ **then**
 - 7: Adv **outputs** $(x_{(L)}^t, x_{(R)}^t) \in \chi^2$
 - 8: $x_{(\text{side})}^t$ is sent to M
 - 9: M **outputs** a_t
-

More formally the relationship between Adv and M is modeled as a game between adversary Adv and mechanism M , given in Game 3.

Definition A.2 (Differential privacy in the adaptive continual release model (Jain et al., 2023)). Given a mechanism M the *view* of the adversary Adv in game $\Pi_{M, Adv}$ (Game 3) consists of Adv 's internal randomness, as well as the outputs of both Adv and M . Let $V_{M, Adv}^{(\text{side})}$ denote Adv 's view at the end of the game run with input side $\in \{L, R\}$. Let \mathcal{V} be the set of all possible views. Mechanism M is (ϵ, δ) -*differentially private in the adaptive continual release model* if, for all adversaries Adv and any $S \subseteq \mathcal{V}$,

$$\Pr(V_{M, Adv}^{(L)} \in S) \leq e^\epsilon \Pr(V_{M, Adv}^{(R)} \in S) + \delta$$

and

$$\Pr(V_{M,Adv}^{(R)} \in S) \leq e^\epsilon \Pr(V_{M,Adv}^{(L)} \in S) + \delta.$$

We also call such a mechanism *adaptively* (ϵ, δ) -differentially private.

Probability Preliminaries

Lemma A.1. *Let Y_1, \dots, Y_k be independent variables with distribution $\text{Lap}(b)$ and let $Y = \sum_{i=1}^k Y_i$. Then*

$$P(|Y| > 2b\sqrt{2\ln(2/\beta_S)} \max(\sqrt{k}, \sqrt{\ln(2/\beta_S)}) \leq \beta_S.$$

Proof. Apply Corollary 12.3 of [Dwork and Roth \(2014\)](#) to $b_1 = \dots = b_k = b$. □

Lemma A.2. *For a random variable $X \sim D$, if $\Pr[|X| > \alpha] \leq \beta$, then for $X_1, X_2, \dots, X_k \sim D$ i.i.d., we have $\Pr[\max_i |X_i| > \alpha] \leq k \cdot \beta$.*

We use $f_X(x)$ to denote the probability density function of a continuous random variable X . For our privacy proofs, we repeatedly use the fact that if X and Y are independent random variables with joint probability density function $f_{X,Y}(x, y)$, then $f_{X,Y}(x, y) = f_X(x) \cdot f_Y(y)$. Thus for any event $A(X, Y)$, we have

$$\int_{x,y} \mathbb{1}[A(x, y)] f_{X,Y}(x, y) dx dy = \int_y \Pr_X[A(X, y)] f_Y(y) dy$$

B HISTOGRAM QUERIES PARAMETERIZED IN MAXIMUM QUERY OUTPUT

We gave an overview of how the mechanism works on an input stream in the main body. Here, we present the privacy and utility proofs of Mechanism 4. We add the variable p_j to the mechanism purely for the proof, to denote the end of the j -th interval. In particular, $[p_{j-1}, p_j]$ is the j -th interval.

B.1 Privacy

Recall that the main technical challenge to prove privacy of Mechanism 4 is the following: The partitioning mechanism (which decides when to close an interval) *depends on the outputs of the histogram mechanism for prior intervals* (unlike in [Dwork et al. \(2015\)](#), where the partitioning was independent of the output of the counting mechanism), and the input to the histogram mechanism depend on the output of the partitioning mechanism, and, hence, on the prior output of the histogram mechanism. Furthermore, given two neighboring streams, the input to the histogram mechanism might not necessarily be neighboring, since the input to the histogram depends on the partitioning (consider the case where two wildly different partitions are used on two neighboring streams, leading to inputs to the histogram mechanism that are very far apart). Thus, we cannot use a simple composition theorem to show privacy for the combined mechanism. To overcome this difficulty, we use a continuous histogram mechanism that is differentially private *even if the inputs are chosen adaptively*. We then perform a careful privacy analysis to show that the interaction between the adaptively differentially private continuous histogram mechanism and the partitioning mechanism satisfies privacy. The fact that the continuous histogram mechanism is adaptively differentially private allows us to separate the privacy loss incurred by the partitioning mechanism from that of the histogram mechanism in the analysis.

Lemma B.1. *Let $\epsilon > 0$. If H is an $(\epsilon/3)$ -differentially private continual histogram mechanism, then Mechanism 4 satisfies ϵ -differential privacy. This holds independent of the initial setting of (s_1, \dots, s_d) , Thresh_k^t , D_j^t , and $C_j^t s$.*

Proof. Let x and y be two neighboring streams that differ at time t^* . Notice that the outputs of Mechanism 4 at any time step are a post-processing of the interval partitioning and the outputs (s_1, \dots, s_d) of the histogram mechanism H for each interval. Thus, to argue privacy, we consider a mechanism $\mathcal{A}(x)$ which outputs the interval partitions and outputs of H for each interval with input stream x . Let S be any subset of possible outputs of $\mathcal{A}()$. We show that

$$\Pr[\mathcal{A}(x) \in S] \leq e^\epsilon \cdot \Pr[\mathcal{A}(y) \in S]$$

Mechanism 4 Mechanism for answering m histogram queries parameterized in the maximum query output.

```

1: Input: Stream  $x^1, x^2, \dots \in \{0, 1\}^d$ , an adaptively  $\epsilon$ -differentially private continual histogram mechanism  $H$ ,
   failure probability  $\beta$ , additive error bound  $\text{err}(t, \beta)$  that holds with probability  $\geq 1 - \beta$  for the output of  $H$ 
   at time step  $t$ .
2: Output: Estimate of  $q_k(h(t))$  for all  $k \in [m]$  and all  $t \in \mathbb{N}$ 
3:  $\triangleright$  Initialization  $\triangleleft$ 
4: Initialize  $H$ 
5:  $\beta' = 6\beta/\pi^2$ ,  $\beta_t = \beta'/t^2$  for any  $t \in \mathbb{N}$ 
6:  $\text{Thresh}_k^1 \leftarrow 3\epsilon^{-1}(12\ln(2/\beta') + 6\ln(6/\beta') + m\ln(6m/\beta')) + 3 \cdot \text{err}(1, \beta'/6)$  for all  $k \in [m]$ 
7:  $c_i \leftarrow 0$  for all  $i \in [d]$   $\triangleright$  column sum within interval
8:  $s_i \leftarrow 0$  for all  $i \in [d]$   $\triangleright$  histogram estimate
9:  $j \leftarrow 1$   $\triangleright$  number of intervals
10:  $p_0 \leftarrow 0$ 
11:  $\tau_1 \leftarrow \text{Lap}(6/\epsilon)$ 
12:  $\text{out} \leftarrow (q_1(\mathbf{0}), q_2(\mathbf{0}), \dots, q_m(\mathbf{0}))$   $\triangleright$  current output
13:  $\triangleright$  Process the input stream  $\triangleleft$ 
14: for  $t \in \mathbb{N}$  do
15:    $c_i \leftarrow c_i + x_i^t$ ,  $s_i \leftarrow s_i + x_i^t$  for all  $i \in [d]$ 
16:    $\triangleright$  Set parameters  $\triangleleft$ 
17:    $\alpha_\mu^t \leftarrow 12\epsilon^{-1}\ln(2/\beta_t)$ 
18:    $\alpha_\tau^j \leftarrow 6\epsilon^{-1}\ln(6/\beta_j)$ 
19:    $\alpha_\gamma^j \leftarrow 3\epsilon^{-1}m\ln(6m/\beta_j)$ 
20:    $\alpha_H^j \leftarrow \text{err}(j, \beta_j/6)$ 
21:    $C_j^t \leftarrow \alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j$ ,  $D_j^t \leftarrow 3(C_j^t + \alpha_H^j)$ 
22:    $\triangleright$  Test for threshold  $\triangleleft$ 
23:    $\mu_t \leftarrow \text{Lap}(12/\epsilon)$ 
24:   if  $\exists k \in [m] : q_k(s) + \mu_t > \text{Thresh}_k^t + \tau_j$  then  $\triangleright$  Close the current interval
25:      $p_j \leftarrow t$ 
26:     insert  $(c_1, \dots, c_d)$  into  $H$ 
27:     reset  $c_i \leftarrow 0$  for all  $i \in [d]$ 
28:     for  $k \in [m]$  do
29:        $\gamma_k^j \leftarrow \text{Lap}(3m/\epsilon)$ 
30:        $\triangleright$  if  $q_k(s)$  is "close" to threshold, increase threshold  $\triangleleft$ 
31:       if  $q_k(s) + \gamma_k^j > \text{Thresh}_k^t - C_j^t$  then
32:          $\text{Thresh}_k^t \leftarrow \text{Thresh}_k^t + D_j^t$ 
33:        $j \leftarrow j + 1$ 
34:        $\text{Thresh}_k^t \leftarrow \text{Thresh}_k^t - D_{j-1}^t + D_j^t$  for all  $k \in [m]$   $\triangleright$  update threshold for the new interval
35:        $\tau_j \leftarrow \text{Lap}(6/\epsilon)$   $\triangleright$  pick fresh noise
36:        $(s_1, \dots, s_d) \leftarrow \text{output}(H)$ 
37:        $\text{out} \leftarrow (q_1(s), \dots, q_m(s))$ 
38:     output  $\text{out}$ 
39:      $\text{Thresh}_k^{t+1} \leftarrow \text{Thresh}_k^t - D_j^t + D_j^{t+1} \forall k \in [m]$ 
40:    $p_j \leftarrow T$ 

```

The arguments also hold when swapping the identities of x and y since they are symmetric, which gives us the privacy guarantee. Thus we focus on proving the inequality above.

We first argue that Mechanism 4 acts like an adversarial process in the adaptive continual release model towards the histogram mechanism H . From our assumption on H it then follows that the output of H is $\epsilon/3$ -differentially private. We will combine this fact with an analysis of the modified sparse vector technique (which determines when to close an interval) plus the properties of the Laplace mechanism (which determines when a threshold is updated) to argue that the combined mechanism consisting of the partitioning and the histogram mechanism is ϵ -differentially private.

Recall that an adversary in the adaptive continual release model presented in Appendix A is given by a privacy

Game 5 Privacy game $\Pi_{H, Adv(x,y)}$ for the adaptive continual release model and m queries for histogram mechanism H

```

1: Input: Streams  $x = x^1, x^2, \dots, x^T \in \{0,1\}^d$  and  $y = y^1, y^2, \dots, y^T \in \{0,1\}^d$  such that  $x$  and  $y$  are
   neighboring and differ in time  $t^*$ , initial values  $s_1, \dots, s_d$ , a stream of values  $D_1, D_2, \dots$ , a stream of values
    $C_1, C_2, \dots$ 
2:  $p_0 \leftarrow 0, j \leftarrow 1$ 
3:  $c_i^x = 0$  and  $c_i^y = 0$  for all  $i \in [d]$ 
4: ChallengeOver = False
5:  $\tau \leftarrow \text{Lap}(6/\epsilon)$ 
6:  $\tilde{D}_{(k)} \leftarrow D_1 + \tau$  for all  $k \in [m]$ 
7: for  $t \in [T]$  do
8:    $c_i^x = c_i^x + x_i^t, s_i = s_i + x_i^t$  for all  $i \in [d]$ 
9:    $c_i^y = c_i^y + y_i^t$  for all  $i \in [d]$ 
10:   $\mu = \text{Lap}(12/\epsilon)$ 
11:  if  $\exists k \in [m]: q_k(s) + \mu > \tilde{D}_{(k)}$  then
12:     $p_j \leftarrow t$ 
13:    if  $p_j \geq t^*$  and ChallengeOver=False then
14:      type $j$  = challenge
15:      output  $(c^x, c^y)$ 
16:      ChallengeOver=True
17:    else
18:      type $j$  = regular
19:      output  $c^x$ 
20:    for  $k \in [m]$  do
21:       $\tilde{q}_k(s) \leftarrow q_k(s) + \text{Lap}(3m/\epsilon)$ 
22:      if  $\tilde{q}_k(s) > D_{(k)} - C_j$  then
23:         $D_{(k)} \leftarrow D_{(k)} + D_j$ 
24:         $j \leftarrow j + 1$ 
25:     $\tau \leftarrow \text{Lap}(6/\epsilon)$ 
26:     $D_{(k)} \leftarrow D_{(k)} - D_{j-1} + D_j$  for all  $k \in [m]$ 
27:     $\tilde{D}_{(k)} \leftarrow D_{(k)} + \tau$  for all  $k \in [m]$ 
28:    reset  $c_i^x \leftarrow 0, c_i^y \leftarrow 0$  for all  $i \in [d]$ 
29:    receive  $(s_1, \dots, s_d) \leftarrow H$ 
30:   $p_j \leftarrow T$ 

```

game, whose generic form is presented in Game 3. Due to the complicated interaction between the partitioning and H , the specification of such an adversarial process in our setting is given in Game 5. Call $[p_{\ell-1}, p_\ell]$ the ℓ -th interval. The basic idea is as follows: Let t^* be the time step at which x and y differ. Conditioned on identical choices for the random variables before time step t^* , we have that all the intervals that the mechanism creates and also the values that the mechanism (in its role as an adversary) gives to the histogram mechanism, are identical for x and y before time step t^* . These are regular time steps in the game. The value for the first interval ending at or after time t^* can differ and constitutes the challenge step. All remaining intervals lead to regular steps in Game 5.

Note that the end of the intervals, i.e., the partitioning of the stream, is computed by the adversary. This partitioning is based on the “noisy” histogram (the s_i values), which are computed from the output of H (which can depend on x and y , depending on side) and the values of the input stream x in the current interval - for *either value* of side, since the adversary does not know side. We denote the adversary with input streams x and y by $Adv(x, y)$, and the corresponding game, Game $\Pi_{H, Adv(x, y)}$. Our discussion above implies that $Adv(x, y)$ does not equal $Adv(y, x)$.

The important observation from this game is that there is only one interval, i.e., only one time step for H , where the adversary outputs two values, and in all other time steps it outputs only one value. Also, at the challenge time step where it sends two values c^x and c^y , these values differ by at most 1. Thus the adversarial process that models the interaction between the partitioning mechanism and H fulfills the condition of the adaptive continual release model. As we assume that H is $\epsilon/3$ -differentially private in that model it follows that for all possible neighboring input streams x and y for $\Pi_{H, Adv(x, y)}$ and all possible sides L and R it holds that

$$\Pr(V_{H, Adv(x, y)}^{(L)} \in S) \leq e^{\epsilon/3} \Pr(V_{H, Adv(x, y)}^{(R)} \in S)$$

where we use the definition of a view $V_{H, Adv(x, y)}^{(L)}$ and $V_{H, Adv(x, y)}^{(R)}$ from Definition A.2. The same also holds with the positions of x and y switched and for L and R switched. Since the choice of L/R merely decides whether the counts c^x or c^y are sent by the game to H , we abuse notation and specify directly which count is sent to H , as $V_{H, Adv(x, y)}^{(x)}$ or $V_{H, Adv(x, y)}^{(y)}$.

Recall that the view of the adversary in Game $\Pi_{H, Adv(x, y)}$ consists of its internal randomness as well as its outputs and the output of H for the whole game, i.e., at the end of the game. The behavior of $Adv(x, y)$ is completely determined by its inputs consisting of x , y , the outputs of H , the thresholds D_j^t and the values C_j^t , as well as by the functions q_k and the random coin flips. However, for the privacy analysis only the partitioning and the output of H matter since the output of Mechanism 4 only depends on those. Thus, we ignore the other values in the view and say that a view V of the adversary $Adv(x, y)$ in Game $\Pi_{H, Adv(x, y)}$ satisfies $V \in S$, if the partitioning and the streams of (s_1, \dots, s_d) returned from H for all intervals match the output sequences in S . Let C_j^t and D_j^t be as in the mechanism. Assume Game $\Pi_{H, Adv(x, y)}$ is run with those settings of C_j^t and D_j^t . By the definition of $\Pi_{H, Adv(x, y)}$, we have

$$\begin{aligned} \Pr(\mathcal{A}(x) \in S) &= \Pr(V_{H, Adv(x, y)}^{(x)} \in S), \text{ and} \\ \Pr(\mathcal{A}(y) \in S) &= \Pr(V_{H, Adv(y, x)}^{(y)} \in S). \end{aligned}$$

We will prove below that

$$\Pr(V_{H, Adv(x, y)}^{(x)} \in S) \leq e^{2\epsilon/3} \Pr(V_{H, Adv(y, x)}^{(x)} \in S). \quad (1)$$

Privacy then follows, since

$$\begin{aligned} \Pr(\mathcal{A}(x) \in S) &= \Pr(V_{H, Adv(x, y)}^{(x)} \in S) \\ &\leq e^{2\epsilon/3} \Pr(V_{H, Adv(y, x)}^{(x)} \in S) \leq e^\epsilon \Pr(V_{H, Adv(y, x)}^{(y)} \in S) \\ &= e^\epsilon \Pr(\mathcal{A}(y) \in S). \end{aligned} \quad (2)$$

We now prove (1). Recall that when we run $Adv(x, y)$ on side x , the interval partitioning is created according to x and the outputs of H . Also for each interval, the input given to H is based on the counts for x , as we consider

side x . When we run $Adv(y, x)$ on side x , then the interval partitioning is created according to y and for each interval we give the counts for x as input to H . Thus in both cases the input given to H is based on the counts for x , and hence, to prove inequality 1, it suffices to show that *when running $Adv(x, y)$ on side x and $Adv(y, x)$ on side x , the probabilities of getting a given partition and thresholds are $e^{2\epsilon/3}$ -close*. To simplify notation, we denote running $Adv(x, y)$ on side x as $run(x)$, and $Adv(y, x)$ on side x as $run(y)$.

Recall that $[p_{\ell-1}, p_\ell]$ is the ℓ^{th} interval. Denote the interval that t^* belongs to as the j -th interval. Note that the probabilities of computing any fixed sequence of intervals $[p_0, p_1), \dots, [p_{j-2}, p_{j-1})$ with $p_{j-1} < t^*$ are the same on both $run(x)$ and $run(y)$, since the streams are equal at all time steps before t^* .

We want to argue two things: (A) fixing a particular time $\lambda > p_{j-1}$, the probability of $p_j = \lambda$ is $e^{\epsilon/3}$ -close on $run(x)$ and $run(y)$; and (B) the probabilities of updating the thresholds, i.e., executing line 23 in Game 5 at time p_j for any subset of $[d]$, is $e^{\epsilon/3}$ -close on $run(x)$ and $run(y)$. Then we show that this implies that (C) all the thresholds $Thresh_k^t$ maintained by adversary are the same at the end of the interval.

The proof of (A) is similar to the privacy of the sparse vector technique (see e.g. [Lyu et al. \(2017\)](#)); (B) holds by a post-processing of the Laplace mechanism; and (C) follows by carefully analyzing the sequences of events and their dependencies. Before we prove these statements, (A), (B) and (C) together imply that the probabilities that the j -th interval ends at the same time *and* that the thresholds are updated in the same way in all intervals in $run(x)$ and $run(y)$ are $e^{2\epsilon/3}$ -close. This implies that the probabilities $\Pr(V_{H, Adv(x, y)}^{(x)} \in S)$ and $\Pr(V_{H, Adv(y, x)}^{(x)} \in S)$ are $e^{2\epsilon/3}$ -close for any subset S of possible outputs. Thus, (1) and therefore (2) follow, completing the proof.

(A) Fixing a particular time $\lambda > p_{j-1}$, we first show that the probability of interval j ending at λ (i.e., $p_j = \lambda$) is $e^{\epsilon/3}$ -close on $run(x)$ and $run(y)$. Fixing some notation, let $\mu_t \sim \text{Lap}(12/\epsilon)$ and $\tau_j \sim \text{Lap}(6/\epsilon)$ be as in the mechanism, let $s^t(x)$ denote the vector of $(s_i)_{i \in [d]}$ at time t for stream x , and f_X denote the density function of the random variable X . For the interval j to close at time λ on $run(x)$, there must exist a $k \in [m]$ with $q_k(s^\lambda(x)) + \mu_\lambda > \text{Thresh}_k^t + \tau_j$ at time λ , and $q_\ell(s^t(x)) + \mu_t \leq \text{Thresh}_\ell^t + \tau_j$ for all $p_{j-1} < t < \lambda$ and $\ell \in [m]$.

Note that conditioning on all the random variables being the same on x and y before p_{j-1} , we have that any s_ℓ at time $t \leq p_j$ can differ by at most 1 on x and y . Therefore $q_\ell(s^t(x))$ and $q_\ell(s^t(y))$ can also differ by at most 1 by sensitivity of q_ℓ . Therefore, for $p_{j-1} < t < \lambda$, any $\ell \in [m]$ and any fixed value $z \in \mathbb{R}$ that τ_j can take, we have

$$\begin{aligned} & \Pr[q_\ell(s^t(x)) + \mu_t \leq \text{Thresh}_\ell^t + z] \\ & \leq \Pr[q_\ell(s^t(y)) + \mu_t \leq \text{Thresh}_\ell^t + z + 1] \end{aligned}$$

Also, for fixed $z \in \mathbb{R}$ (resp. $c \in \mathbb{R}$) that τ_j (resp. μ_λ) can take,

$$\begin{aligned} & \Pr[q_k(s^\lambda(x)) + c > \text{Thresh}_k^t + z] \\ & \leq \Pr[q_k(s^\lambda(y)) + c + 2 > \text{Thresh}_k^t + z + 1]. \end{aligned}$$

Now, since $\tau_j \sim \text{Lap}(6/\epsilon)$, we have $f_{\tau_j}(z) \leq e^{\epsilon/6} f_{\tau_j}(z + 1)$. Similarly, since $\mu_\lambda \sim \text{Lap}(12/\epsilon)$, we have $f_{\mu_\lambda}(c) \leq e^{2\epsilon/12} f_{\mu_\lambda}(c + 2) = e^{\epsilon/6} f_{\mu_\lambda}(c)$. Now, integrating over the distributions of τ_j and μ_λ and using these properties gives $\Pr[p_j = \lambda \text{ on } x] \leq e^{\epsilon/3} \Pr[p_j = \lambda \text{ on } y]$. We conclude that the probability of $p_j = \lambda$ is $e^{\epsilon/3}$ -close on $run(x)$ and $run(y)$.

(B) Next, conditioned on all previous outputs of H being the same and p_j being equal, we argue that the probabilities of updating any subset of thresholds are close for both runs at time p_j . Note that when they are updated at the same time, they are updated in the same way. Since $q_k(s^{p_j}(x))$ and $q_k(s^{p_j}(y))$ can differ by at most 1 for each $k \in [m]$, adding $\gamma_k^j \sim \text{Lap}(3m/\epsilon)$ to every $q_k(s^{p_j}(y))$ in line 22 ensures that the distributions of $q_k(s^{p_j}(x)) + \gamma_k^j$ and $q_k(s^{p_j}(y)) + \gamma_k^j$ are $e^{\epsilon/3}$ -close for all $k \in [m]$ by composition. Since the condition in line 23 only depends on those, this implies that the probabilities of updating the threshold (i.e., executing line 23) on any subset of $[m]$ on $run(x)$ and $run(y)$ are $e^{\epsilon/3}$ -close.

(C) *Up to interval $j - 1$:* We already argued in (A) that conditioned on all random variables being the same on x and y before interval j , the executions of $run(x)$ and $run(y)$ are identical and, thus, all thresholds are updated in the same way. *Interval j and up:* For any $\ell \geq j$ denote by E_ℓ the event that for $run(x)$ and $run(y)$, all the intervals until interval ℓ end at the same time step, all the thresholds $Thresh_k^t$ for $t \leq p_\ell$ are identical, and the random variables used after time p_ℓ take the same values on both runs. We will next argue that conditioned on event E_ℓ , event $E_{\ell+1}$ holds. Note that event E_j holds by (B), and by definition, $run(x)$ and $run(y)$ both use the

counts from stream x to compute the input for H . Inductively assume that event E_ℓ holds. Event E_ℓ implies that all intervals $\leq \ell$ were closed at the same time on both runs and hence the same counts were given as input to H . Since (a) the streams x and y are identical for all $t > p_\ell$, (b) the thresholds and the outputs of H are identical at the end of interval ℓ , and (c) the random variables used after p_ℓ are identical (which follows from event E_ℓ), we have that the $\ell + 1$ -st interval ends at the same time on both runs, and that the same thresholds are updated, and by the same amount at time $p_{\ell+1}$. This shows that event $E_{\ell+1}$ holds, as required. \square

B.2 Accuracy

After processing the input at time step t , let h^t be the actual histogram, s^t be the value of s stored by Mechanism 1, and q^{*t} be the maximum query value. Suppose t belongs to interval j , i.e., $t \in [p_{j-1}, p_j]$. Since the mechanism outputs $q_k(s^{p_{j-1}})$ at time t , our goal is to bound the additive error $|q_k(h^t) - q_k(s^{p_{j-1}})|$ at all times $t \in \mathbb{N}$ and for all queries $k \in [m]$. We do this as follows:

1. Use Laplace concentration bounds to bound the maximum value attained by the random variables used by the mechanism (Lemma B.2).
2. Show that if query k crosses the threshold Thresh_k^t , then q_k on the true histogram is not too much smaller than the threshold (Lemma B.4).
3. Show that if query k crosses the threshold Thresh_k^t , then q_k on the true histogram is not too much larger than the threshold (Lemma B.5).
4. Bound the number of intervals produced by the mechanism (Lemma B.6).
5. Use all the above to bound the error of the mechanism (Lemma B.7).

We define the random variables (RVs) $\mu_t, \tau_j, \gamma_k^j$ as in the mechanism. The variables $\alpha_\mu^t, \alpha_\tau^j, \alpha_\gamma^j$ used in the mechanism are defined such that they bound simultaneously with good probability ($\geq 1 - \beta$) the corresponding RVs. In the rest of the section, we condition that the bounds hold on the random variables used in the mechanism.

Lemma B.2 (RV Bounds). *There exists a histogram mechanism H such that the following bounds hold simultaneously with probability $\geq 1 - \beta$ for all $t, j \in \mathbb{N}$ and $k \in [m]$*

$$|\mu_t| \leq \alpha_\mu^t, \quad |\tau_j| \leq \alpha_\tau^j, \quad |\gamma_k^j| \leq \alpha_\gamma^j, \quad \max_{t \in [p_{j-1}, p_j]} \|s^t - h^t\|_\infty \leq \alpha_H^j \quad \forall t \in [p_{j-1}, p_j]$$

where

$$\alpha_\mu^t = 12\epsilon^{-1} \ln(2/\beta_t), \quad \alpha_\tau^j = 6\epsilon^{-1} \ln(6/\beta_j), \quad \alpha_\gamma^j = 3\epsilon^{-1} m \ln(6m/\beta_j),$$

$$\alpha_H^j = O(\epsilon^{-1} d \cdot (\log(j) \log(d/\beta) + (\log j)^{1.5} \sqrt{\log(d/\beta)}))$$

From the final bound above, we get the following lemma which bounds the error of the query values when computed on the noisy histogram s stored by the mechanism.

Lemma B.3. *Assume Lemma B.2 holds. Let $t \in [T]$ be any time step, and suppose $t \in [p_{j-1}, p_j]$. Then for all $k \in [m]$,*

$$|q_k(s^t) - q_k(h^t)| \leq \alpha_H^j.$$

Since our output at time t is $q_k(s^{p_{j-1}})$, our error is $|q_k(h^t) - q_k(s^{p_{j-1}})|$, which we bound as follows:

$$\begin{aligned} |q_k(h^t) - q_k(s^{p_{j-1}})| &\leq |q_k(h^t) - q_k(h^{p_{j-1}})| + |q_k(h^{p_{j-1}}) - q_k(s^{p_{j-1}})| \\ &\leq |q_k(h^t) - q_k(h^{p_{j-1}})| + \alpha_H^j && \text{(by Lemma B.3)} \\ &\leq q_k(h^t) - q_k(h^{p_{j-1}}) + \alpha_H^j, && \text{(since } q_k \text{ and } h \text{ are monotone and } t \geq p_{j-1}) \end{aligned}$$

our accuracy bound reduces to giving an upper bound on $q_k(h^t)$ and a lower bound on $q_k(h^{p_{j-1}})$.

We say k crosses the threshold at time t if line 32 of the mechanism is executed for k at time t . Note that then $t = p_j$ for some j . Our lower bound on $q_k(h^{p_j})$ will be based on the fact that k crosses the threshold at time p_j . At time steps where k did not cross the threshold, our upper bound on $q_k(h^t)$ will follow from a complementary argument to the above lower bound.

For an upper bound on $q_k(h^{p_j})$ at time steps when k crosses the threshold, we first show that k did not cross the threshold at time $p_j - 1$ as follows: Let $p_\ell < p_j$ be the last time step before p_j when k crosses the threshold, and never in between p_ℓ and p_j . Then by definition of the mechanism, $\text{Thresh}_k^{p_j} - \text{Thresh}_k^{p_\ell} = D_j^{p_j}$. We use this to show that q_k must have increased by more than 1 between p_ℓ and p_j . The latter fact implies two things: first, that $j \leq mq^*$; second, that k did not cross the threshold at time $p_j - 1$. The latter can be used to get an upper bound on $q_k(h^{p_j-1})$ and, by the 1-sensitivity of q_k , also on $q_k(h^{p_j})$. For the first interval, there does not exist any such p_ℓ where the threshold was crossed previously. For this, we prove an auxiliary lemma that says that $p_1 > 1$, and hence no threshold was crossed at time $p_1 - 1$, and the rest of the analysis follows.

Combining the two gives an upper bound on $q_k(h^t) - q_k(h^{p_j-1})$ of $O(D_j^t + \alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j)$, which is the crucial bound needed to upper bound $|q_k(h^t) - q_k(h^{p_j-1})|$.

Our first lemma shows that whenever k crosses the threshold, the query value on the true histogram is not too small compared to the threshold.

Lemma B.4 (lower bound). *Assume Lemma B.2 holds. Let $k \in [m]$ and suppose k crosses the threshold at time $t = p_j$.*

$$q_k(h^{p_j}) \geq \text{Thresh}_k^{p_j} - \left(\alpha_\mu^{p_j} + \alpha_\tau^j + 2\alpha_\gamma^j + \alpha_H^j \right).$$

Using the strategy mentioned above, we then prove that the query value on the true histogram is never too large compared to the threshold. Along the way, we also show that every time k crosses the threshold, the query value on the true histogram must increase.

Lemma B.5 (upper bound). *Assume Lemma B.2 holds. Let $k \in [m]$ and $t \in \mathbb{N}$.*

$$q_k(h^t) < \text{Thresh}_k^t + \left(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j + 1 \right).$$

Further, suppose k crosses the threshold at time $t = p_j$. Then denoting by p_ℓ the last time before p_j that k crossed a threshold, it also holds that $p_j - p_\ell > 1$ and $|q_k(h^{p_j}) - q_k(h^{p_\ell})| > 1$.

We use the second part of the above lemma to bound the number of intervals created by the mechanism, where q^{*t} is the maximum query output at time t .

Lemma B.6. *Assume Lemma B.2 holds. Mechanism 1 creates at most mq^{*t} many segments upto time t .*

Lemma B.7. *Assume Lemma B.2 holds. Let $t \in \mathbb{N}$ be any time step, and suppose $t \in [p_{j-1}, p_j]$. Then Mechanism 1 is $\alpha_j^{(t)}$ -accurate at time t , where*

$$\alpha_j^{(t)} = O\left(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j\right)$$

In particular, for all $t \in \mathbb{N}$, Mechanism 1 is $\alpha^{(t)}$ -accurate, where

$$\alpha^{(t)} = O\left(\alpha_\mu^t + \alpha_\tau^{mq^{*t}} + \alpha_\gamma^{mq^{*t}} + \alpha_H^{mq^{*t}}\right).$$

The accuracy proof for Mechanism 4 then follows since we show that Lemma B.2 holds for the corresponding values in the mechanism, and plugging them into the above lemma.

Corollary B.8. *Mechanism 4 with a histogram mechanism with error $\text{err}(t, \beta)$ has error at most*

$$\alpha^{(t)} = O\left(\frac{1}{\epsilon}(d \text{err}(mq^{*t}, \beta/(mq^{*t})^2) + m \log(mq^{*t}/\beta) + \log t)\right)$$

at all time steps t simultaneously with probability at least $1 - \beta$. In particular, using the histogram mechanism from Fact 5 has error at most

$$\alpha^{(t)} = O\left((d \log^2(dm q^*/\beta) + m \log(mq^*/\beta) + \log t) \epsilon^{-1}\right)$$

at all time steps t simultaneously with probability at least $1 - \beta$.

B.3 Accuracy Proofs

Lemma B.2 (RV Bounds). *There exists a histogram mechanism H such that the following bounds hold simultaneously with probability $\geq 1 - \beta$ for all $t, j \in \mathbb{N}$ and $k \in [m]$*

$$\begin{aligned} |\mu_t| &\leq \alpha_\mu^t, & |\tau_j| &\leq \alpha_\tau^j, & |\gamma_k^j| &\leq \alpha_\gamma^j, & \max_{t \in [p_{j-1}, p_j)} \|s^t - h^t\|_\infty &\leq \alpha_H^j \quad \forall t \in [p_{j-1}, p_j) \\ \text{where} \quad \alpha_\mu^t &= 12\epsilon^{-1} \ln(2/\beta_t), & \alpha_\tau^j &= 6\epsilon^{-1} \ln(6/\beta_j), & \alpha_\gamma^j &= 3\epsilon^{-1} m \ln(6m/\beta_j), \\ \alpha_H^j &= O(\epsilon^{-1} d \cdot (\log(j) \log(d/\beta) + (\log j)^{1.5} \sqrt{\log(d/\beta)})) \end{aligned}$$

Proof. Using Lemma 2 gives us the first three bounds below:

1. $\mu_t \sim \text{Lap}(12/\epsilon)$ satisfies $|\mu_t| < 12\epsilon^{-1} \ln(2/\beta_t)$ with probability $\geq 1 - \beta_t/2$.
2. $\tau_j \sim \text{Lap}(6/\epsilon)$ satisfies $|\tau_j| < 6\epsilon^{-1} \ln(6/\beta_j)$ with probability $\geq 1 - \beta_j/6$.
3. $\gamma_k^j \sim \text{Lap}(3m/\epsilon)$ satisfies $|\gamma_k^j| \leq 3\epsilon^{-1} m \ln(6m/\beta_j)$ for all $k \in [m]$ simultaneously with probability $\geq 1 - \beta_j/6$.
4. By assumption, the output of H at time p_{j-1} has additive error at most $\text{err}(j, \beta_j/6)$ with probability at least $1 - \beta_j/6$. In particular, the histogram mechanism from Fact 5 guarantees $\text{err}(j, \beta) = O(\epsilon^{-1} d \cdot (\log(j) \log(d/\beta) + (\log j)^{1.5} \sqrt{\log(d/\beta)}))$.

By a union bound, all the four bounds hold at every time step with probability at least $1 - \sum_{t=1}^{\infty} \beta_t/2 - \sum_{j=1}^{\infty} \beta_j/6 = 1 - \beta$. \square

Lemma B.3. *Assume Lemma B.2 holds. Let $t \in [T]$ be any time step, and suppose $t \in [p_{j-1}, p_j)$. Then for all $k \in [m]$,*

$$|q_k(s^t) - q_k(h^t)| \leq \alpha_H^j.$$

Proof. This follows, since

$$|q_k(s^t) - q_k(h^t)| \leq \|s^t - h^t\|_\infty \leq \alpha_H^j$$

where the first inequality is a consequence of q_k having sensitivity one, and the second is from the Lemma B.2. \square

Lemma B.4 (lower bound). *Assume Lemma B.2 holds. Let $k \in [m]$ and suppose k crosses the threshold at time $t = p_j$.*

$$q_k(h^{p_j}) \geq \text{Thresh}_k^{p_j} - \left(\alpha_\mu^{p_j} + \alpha_\tau^j + 2\alpha_\gamma^j + \alpha_H^j \right).$$

Proof. This follows from the sensitivity of q_k and the fact that k crosses the threshold at time p_j .

$$\begin{aligned} q_k(h^{p_j}) &\geq q_k(s^{p_j}) - \alpha_H^j && \text{(by Lemma B.3)} \\ &\geq q_k(s^{p_j}) + \gamma_k^j - \alpha_\gamma^j - \alpha_H^j && \text{(by definition of } \alpha_\gamma^j) \\ &\geq \text{Thresh}_k^{p_j} - C_j^{p_j} - \alpha_\gamma^j - \alpha_H^j && \text{(since } k \text{ crosses the threshold)} \\ &\geq \text{Thresh}_k^{p_j} - \alpha_\mu^t - \alpha_\tau^j - 2\alpha_\gamma^j - \alpha_H^j && \text{(by definition of } C_j^{p_j}) \end{aligned}$$

as required. \square

Lemma B.9. *Assume Lemma B.2 holds. Let $k \in [m]$ and suppose k did not cross the threshold at time t . Then*

$$q_k(h^t) < \text{Thresh}_k^t + \left(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j \right).$$

Proof. Since k did not cross the threshold at time t , either the condition in line 24 was false or the condition in line 31 was false for k at time t . Thus, one of the following holds

$$\begin{aligned} q_k(s^t) &< \text{Thresh}_k^t + \alpha_\mu^t + \alpha_\tau^j < \text{Thresh}_k^t + \alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j && \text{if line 24 was false, or} \\ q_k(s^t) &< \text{Thresh}_k^t - C_j^t + \alpha_\gamma^j < \text{Thresh}_k^t + \alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j && \text{if line 31 was false.} \end{aligned}$$

Combining this with Lemma B.3 gives the required bound. \square

Lemma B.10. *Assume Lemma B.2 holds. No interval is closed on the first time step, i.e., $p_1 > 1$.*

Proof. Note that $C_j^t = \alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j$. Thus, if the condition in line 24 is true at time p_j , then the condition in line 31 is also true for some k . Said differently, whenever we end a segment, there also exists an k such that k crosses the threshold. Using Lemma B.4 with $t = p_1$ gives us that

$$q_k(h^{p_1}) \geq \text{Thresh}_k^{p_1} - (\alpha_\mu^{p_1} + \alpha_\tau^1 + 2\alpha_\gamma^1 + \alpha_H^1).$$

Note that since $D_1^{p_1} > \alpha_\mu^{p_1} + \alpha_\tau^1 + 2\alpha_\gamma^1 + \alpha_H^1$, this implies $q_k(h^{p_1}) > 1$. As q_k increases by at most 1 per time step and $q_k(0, \dots, 0) = 0$, it follows that $p_1 > 1$. \square

Lemma B.5 (upper bound). *Assume Lemma B.2 holds. Let $k \in [m]$ and $t \in \mathbb{N}$.*

$$q_k(h^t) < \text{Thresh}_k^t + (\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j + 1).$$

Further, suppose k crosses the threshold at time $t = p_j$. Then denoting by p_ℓ the last time before p_j that k crossed a threshold, it also holds that $p_j - p_\ell > 1$ and $|q_k(h^{p_j}) - q_k(h^{p_\ell})| > 1$.

Proof. If k did not cross the threshold at time t , then the bound follows from Lemma B.9. Thus assume k crosses the threshold at time $t = p_j$. The first part of the claim follows if we show that k did not cross the threshold at time $p_j - 1$, and $p_j - 1 \geq 1$, since then Lemma B.9 holds at time $p_j - 1$ and q_k has sensitivity one. We show the claim by induction over the number of times k crosses the threshold.

Case 1: p_j is the first time k crosses the threshold. Since p_j is the first time k crosses the threshold, clearly, k did not cross the threshold at time $p_j - 1$. Further, Lemma B.10 gives us that $p_j \geq p_1 > 1$ and therefore $p_j - 1 \geq 1$. Using Lemma B.9 with $t = p_j - 1$, and the fact that q_k has sensitivity one gives the required bound.

Case 2: p_j is not the first time k crosses the threshold. Clearly $p_j - 1 \geq 1$ holds in this case. Then let p_ℓ be the last time at which k crosses the threshold before p_j . By induction, we have for p_ℓ that

$$\begin{aligned} q_k(h^{p_\ell}) &< \text{Thresh}_k^{p_\ell} + \alpha_\mu^{p_\ell} + \alpha_\tau^\ell + \alpha_\gamma^\ell + \alpha_H^\ell + 1 \\ &\leq \text{Thresh}_k^{p_j} - D_j^{p_j} + \alpha_\mu^{p_j} + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j + 1 \end{aligned}$$

Since k crosses the threshold at time p_j , Lemma B.4 with $t = p_j$ gives

$$q_k(h^{p_j}) \geq \text{Thresh}_k^{p_j} - (\alpha_\mu^{p_j} + \alpha_\tau^j + 2\alpha_\gamma^j + \alpha_H^j)$$

Putting both these inequalities together, we get

$$\begin{aligned} |q_k(h^{p_j}) - q_k(h^{p_\ell})| &> \left(\text{Thresh}_k^{p_j} - (\alpha_\mu^{p_j} + \alpha_\tau^j + 2\alpha_\gamma^j + \alpha_H^j) \right) \\ &\quad - \left(\text{Thresh}_k^{p_j} - D_j^{p_j} + (\alpha_\mu^{p_j} + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j + 1) \right) \\ &= D_j^{p_j} - (2\alpha_\mu^{p_j} + 2\alpha_\tau^j + 3\alpha_\gamma^j + 2\alpha_H^j + 1) > 1, \end{aligned}$$

since $D_j^t \geq 3(C_j^t + \alpha_H^j)$ and $C_j^t = \alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j$. As q_k has sensitivity one, we have $p_j - p_\ell > 1$, and thus, k did not cross the threshold at time $p_j - 1$. Lemma B.9 with $t = p_j - 1$ and the sensitivity of q_k then gives the required bound. \square

Lemma B.6. Assume Lemma B.2 holds. Mechanism 1 creates at most mq^{*t} many segments upto time t .

Proof. Since $C_j^t = \alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j$, whenever the condition in line 24 is true, then the condition in line 31 is also true for some k , i.e., k crosses the threshold. By Lemma B.5, the query value of q_k on the true histogram grows by at least one every time k crosses the threshold. Since q^{*t} bounds the maximum number of times any query answer can increase before time t , there can be at most mq^{*t} many threshold crossings for all $k \in [m]$ combined, and thus the lemma follows. \square

Lemma B.7. Assume Lemma B.2 holds. Let $t \in \mathbb{N}$ be any time step, and suppose $t \in [p_{j-1}, p_j)$. Then Mechanism 1 is $\alpha_j^{(t)}$ -accurate at time t , where

$$\alpha_j^{(t)} = O\left(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j\right)$$

In particular, for all $t \in \mathbb{N}$, Mechanism 1 is $\alpha^{(t)}$ -accurate, where

$$\alpha^{(t)} = O\left(\alpha_\mu^t + \alpha_\tau^{mq^{*t}} + \alpha_\gamma^{mq^{*t}} + \alpha_H^{mq^{*t}}\right).$$

Proof. Once we prove the first part, the second follows from Lemma B.6. Since $t \in [p_{j-1}, p_j)$, the output of the mechanism at time t is $q_k(s^{p_{j-1}})$. Thus the error at time t is

$$\begin{aligned} |q_k(h^t) - q_k(s^{p_{j-1}})| &\leq |q_k(h^t) - q_k(h^{p_{j-1}})| + |q_k(h^{p_{j-1}}) - q_k(s^{p_{j-1}})| \\ &\leq |q_k(h^t) - q_k(h^{p_{j-1}})| + \alpha_H^j && \text{(by Lemma B.3)} \\ &\leq q_k(h^t) - q_k(h^{p_{j-1}}) + \alpha_H^j && \text{(since } q_k \text{ monotone and } t \geq p_{j-1}) \end{aligned}$$

Our task reduces to giving an upper bound on $q_k(h^t)$, and a lower bound on $q_k(h^{p_{j-1}})$. We have two cases depending on whether k has previously crossed a threshold. Let $t_{\text{first}}(k)$ be the first time in the whole input sequence that k crosses the threshold.

Case 1: $t < t_{\text{first}}(k)$. Since the histogram is empty before the first input arrives, $q_k(h^{p_0}) = 0$. Thus

$$\begin{aligned} q_k(h^t) - q_k(h^{p_{j-1}}) &\leq q_k(h^t) - q_k(h^{p_0}) \\ &< \text{Thresh}_k^t + (\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j + 1) && \text{(by Lemma B.5)} \\ &= D_j^t + (\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j + 1) && \text{(since } \text{Thresh}_k^t = D_j^t) \\ &= O(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j) && \text{(since } D_j^t = 3(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j)) \end{aligned}$$

Case 2: $t \geq t_{\text{first}}(k)$. Let p_ℓ be the largest time step before t when k crosses the threshold. Then

$$\begin{aligned} q_k(h^t) &\leq \text{Thresh}_k^t + (\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j + 1) && \text{(by Lemma B.5)} \\ \text{and } q_k(h^{p_\ell}) &\geq \text{Thresh}_k^{p_\ell} - (\alpha_\mu^{p_\ell} + \alpha_\tau^\ell + 2\alpha_\gamma^\ell + \alpha_H^\ell) && \text{(by Lemma B.4)} \\ &\geq \text{Thresh}_k^t - D_j^t - (\alpha_\mu^t + \alpha_\tau^j + 2\alpha_\gamma^j + \alpha_H^j) \end{aligned}$$

Putting these together, we get

$$\begin{aligned} q_k(h^t) - q_k(h^{p_{j-1}}) &\leq q_k(h^t) - q_k(h^{p_\ell}) \\ &= O(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j) && \text{(since } D_j^t = 3(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j)) \end{aligned}$$

which proves the lemma. \square

Corollary B.8. Mechanism 4 with a histogram mechanism with error $\text{err}(t, \beta)$ has error at most

$$\alpha^{(t)} = O\left(\frac{1}{\epsilon}(d\text{err}(mq^{*t}, \beta/(mq^{*t})^2) + m \log(mq^{*t}/\beta) + \log t)\right)$$

at all time steps t simultaneously with probability at least $1 - \beta$. In particular, using the histogram mechanism from Fact 5 has error at most

$$\alpha^{(t)} = O\left((d \log^2(dm q^*/\beta) + m \log(m q^*/\beta) + \log t) \epsilon^{-1}\right)$$

at all time steps t simultaneously with probability at least $1 - \beta$.

Proof. Lemma B.2 and Lemma B.7 together give us that Mechanism 4 is $\alpha^{(t)}$ -accurate at time t , where

$$\begin{aligned} \alpha_\mu^t &= O\left(\epsilon^{-1} \log(2t/\beta)\right), & \alpha_\tau^j &= O\left(\epsilon^{-1} \log(j/\beta)\right), \\ \alpha_\gamma^j &= O\left(\epsilon^{-1} m \log(mj/\beta)\right), & \alpha_H^j &= O(\text{err}(j, \beta/j^2)) \end{aligned}$$

Since $\alpha^{(t)} = O\left(\alpha_\mu^t + \alpha_\tau^j + \alpha_\gamma^j + \alpha_H^j\right)$ with $j \leq m q^{*t}$, this gives the lemma. For the histogram mechanism from Fact 5,

$$\begin{aligned} \alpha_H^j &= O\left(\epsilon^{-1} d \cdot \left(\log(j) \log(dj/\beta) + (\log j)^{1.5} \sqrt{\log(dj/\beta)}\right)\right) \\ &= O\left(\epsilon^{-1} d \log^2(dj/\beta)\right), \end{aligned}$$

which gives

$$\alpha^{(t)} = O\left((d \log^2(dm q^*/\beta) + m \log(m q^*/\beta) + \log t) \epsilon^{-1}\right)$$

as claimed. \square

B.4 Extensions

For (ϵ, δ) -dp, we use an adaptively differentially private continual histogram mechanism H and the Gaussian mechanism for γ_k^j , which gives an error bound of

$$\alpha^{(t)} = O\left(\epsilon^{-1} \log(1/\delta) \cdot \left(\sqrt{d} \log^{1.5}(dm q^{*t}/\beta) + \sqrt{m} \log(m q^{*t}/\beta) + \log t\right)\right)$$

for (ϵ, δ) -differential privacy. We present the technical details in Appendix F.

C HISTOGRAM PARAMETERIZED IN THE NUMBER OF FLUCTUATIONS

As earlier, we gave an overview of how the mechanism works on an input stream in the main body. Here, we present the privacy and utility proofs of Mechanism 6. We add the variable p_j to the mechanism purely for the proof, to denote the end of the j -th interval. In particular, $[p_{j-1}, p_j)$ is the j -th interval.

C.1 Privacy

We show ϵ -differential privacy for neighboring streams x and y which are allowed to differ at one time step by ℓ_∞ -norm 1, i.e., there exists a t^* such that $\|x^{t^*} - y^{t^*}\|_\infty \leq 1$. Then, by group privacy, the mechanism is 2ϵ -differentially private for neighboring streams that are allowed to differ by ℓ_∞ -norm 2.

Lemma C.1. *Mechanism 6 is ϵ -differentially private.*

To prove Lemma C.1, we note that the outputs of Mechanism 6 are a post-processing of three parts: 1.) the mechanism computing p_0, p_1, \dots ; 2.) the mechanism updating the values of *mode*; 3.) and a continual histogram mechanism. Note that in Mechanism 6, 1.) and 2.) do *not* depend on the outputs of 3.), and 3.) is $\epsilon/3$ -differentially private by assumption. Thus, it is enough to show that the parts of the mechanism computing p_0, p_1, \dots and the values of *mode* together are $2\epsilon/3$ -differentially private, as then Mechanism 6 is differentially private by the composition theorem (Fact 3).

Lemma C.2. *The mechanism obtained by running Mechanism 6 and outputting only the values of p_0, p_1, \dots , and the values of mode_i at every time step, is $2\epsilon/3$ -differentially private.*

Mechanism 6 Mechanism for HISTOGRAM parameterized in the number of fluctuations.

```

1: Input: Stream  $x^1, x^2, \dots \in \{-1, 0, 1\}^d$ , an  $\epsilon/3$ -differentially private continual histogram mechanism  $H$ ,
   failure probability  $\beta$ , additive error bound  $\text{err}(t, \beta)$  that holds with probability  $\geq 1 - \beta$  for the output of  $H$ 
   at time step  $t$ .
2: Output: Estimate  $h(t)$  at all  $t \in \mathbb{N}$ 
3:  $\triangleright$  Initialization of all parameters  $\triangleleft$ 
4: Initialize  $H$ 
5:  $\beta' = 6\beta/\pi^2$ ,  $\beta_t = \beta'/t^2$  for any  $t \in \mathbb{N}$ 
6:  $j \leftarrow 1$   $\triangleright$  number of intervals
7:  $p_0 \leftarrow 0$ 
8:  $c_i \leftarrow 0$  for all  $i \in [d]$   $\triangleright$  column sum within interval
9:  $\text{mode}_i \leftarrow 0$  for all  $i \in [d]$ 
10:  $t_{\text{diff}} \leftarrow 0$   $\triangleright$  length of current interval
11:  $\tau_1 \leftarrow \text{Lap}(9/\epsilon)$ ,  $\tau_2 \leftarrow \text{Lap}(9/\epsilon)$ 
12:  $H_{\text{out}} = 0^d$   $\triangleright$  initial histogram
13:  $\triangleright$  Process the input stream  $\triangleleft$ 
14: for  $t \in \mathbb{N}$  do
15:    $c_i \leftarrow c_i + x_i^t$  for all  $i \in [d]$ 
16:    $t_{\text{diff}} \leftarrow t_{\text{diff}} + 1$ 
17:    $\alpha_t \leftarrow \frac{27}{\epsilon} \log(4/\beta_t) + \frac{3d}{\epsilon} \log(1/\beta_j)$ 
18:    $\text{Thresh}_{i,1} \leftarrow \text{mode}_i \cdot t_{\text{diff}} - 2\alpha_t$ ,  $i \in [d]$ 
19:    $\text{Thresh}_{i,2} \leftarrow \text{mode}_i \cdot t_{\text{diff}} + 2\alpha_t$ ,  $i \in [d]$ 
20:    $\mu_1^t \leftarrow \text{Lap}(18/\epsilon)$ 
21:    $\mu_2^t \leftarrow \text{Lap}(18/\epsilon)$ 
22:   if  $\min_i(c_i - \text{Thresh}_{i,1}) < \tau_1 - \mu_1^t$  then
23:      $p_j \leftarrow t$   $\triangleright$  close the current interval
24:     insert  $(c_1, \dots, c_d)$  into  $H$ 
25:      $H_{\text{out}} \leftarrow \text{output}(H)$   $\triangleright$  update histogram
26:     for  $i \in [d]$  do
27:        $\lambda_i = \text{Lap}(3d/\epsilon)$ 
28:        $\triangleright$  update modes  $\triangleleft$ 
29:       if  $c_i + \lambda_i < \text{Thresh}_{i,1} + \alpha_t$  then
30:          $\text{mode}_i \leftarrow \max\{\text{mode}_i - 1, -1\}$ 
31:       reset  $c_i \leftarrow 0$  for all  $i \in [d]$ 
32:        $j \leftarrow j + 1$ 
33:        $t_{\text{diff}} \leftarrow 0$ ;  $\tau_1 \leftarrow \text{Lap}(9/\epsilon)$ 
34:   else if  $\max_i(c_i - \text{Thresh}_{i,2}) > \tau_2 - \mu_2^t$  then
35:      $p_j \leftarrow t$   $\triangleright$  close the current interval
36:     insert  $(c_1, \dots, c_d)$  into  $H$ 
37:      $H_{\text{out}} \leftarrow \text{output}(H)$   $\triangleright$  update histogram
38:     for  $i \in [d]$  do
39:        $\lambda_i = \text{Lap}(3d/\epsilon)$ 
40:        $\triangleright$  update modes  $\triangleleft$ 
41:       if  $c_i + \lambda_i > \text{Thresh}_{i,2} - \alpha_t$  then
42:          $\text{mode}_i \leftarrow \min\{\text{mode}_i + 1, 1\}$ 
43:       reset  $c_i \leftarrow 0$  for all  $i \in [d]$ 
44:        $j \leftarrow j + 1$ 
45:        $t_{\text{diff}} \leftarrow 0$ ;  $\tau_2 \leftarrow \text{Lap}(9/\epsilon)$ 
46:   output  $H_{\text{out}} + \text{mode} \cdot t_{\text{diff}}$ 
47:  $p_j \leftarrow T$ 

```

We first provide a short proof sketch. The partitioning mechanism basically consists of instances of AboveThreshold (see [Dwork and Roth \(2014\)](#)) on disjoint parts of the stream, and the updating of modes is a post-processing of a Laplace mechanism. Thus, intuitively, it should be enough to use parallel composition on both the AboveThreshold mechanisms and the Laplace mechanisms, and then use sequential composition on 1.) and 2.). However, there is a technicality that the inputs and thresholds to AboveThreshold mechanisms depend on their previous outputs, and the mode updates depend on the output of the AboveThreshold mechanism, which depends on the previous mode update. Thus, we cannot use a regular parallel composition argument and thus provide the full proof.

Proof. We focus on proving that the part of Mechanism 6 computing p_0, p_1, \dots and the sequence of *mode* values is $2\epsilon/3$ -differentially private. Since these do not depend on the outputs of 3.) the histogram mechanism, we can then use composition to argue that Mechanism 6 is ϵ -differentially private under continual observation.

Consider any possible partitioning $P = [p_0, p_1), \dots, [p_{\ell-1}, p_\ell)$ of $[0, T)$ and any stream of $M = M^0, \dots, M^T$, where $M^t = (\text{mode}_1^t, \dots, \text{mode}_d^t)$ and mode_i^t is the setting of variable mode_i at time t . We show that the probabilities of getting P and M when running Mechanism 6 on two neighboring streams are $2\epsilon/3$ -close. For two neighboring streams x and y , let t^* be the time where x and y differ and let $[p_{j-1}, p_j)$ be the interval in P which contains t^* . Note that $\|x^{t^*} - y^{t^*}\|_\infty \leq 1$. Before p_{j-1} , all probabilities are the same; conditioning on the run of x and y to be identical up to time p_{j-1} , we show that the probabilities of closing the next interval of y at p_j and updating the mode of y to $\text{mode}_i^{p_j}$ are close to the probabilities of doing so on x . Note that the original setting of mode_i is always 0 for all i on both x and y . Let $c_i^t(x)$ and $c_i^t(y)$ denote the values of c_i at time t on the run of x and y , respectively.

We compare the probabilities of closing the j th interval at time p_j for x and y , and doing so because the conditions in line 22 resp. 34 were fulfilled (note that we need to differentiate these two, since the values of *mode* get updated differently in both cases). In order to close the j -th interval at time p_j , either the condition in line 22 or 34 has to be fulfilled at time p_j , and neither of them can be fulfilled at any time $t \in (p_{j-1}, p_j)$. First, we analyze the probabilities that at time p_j , condition 22 was fulfilled, and neither condition 34 nor 22 was fulfilled at times $t \in (p_{j-1}, p_j)$ (case A). Let z_1 be some fixed value for τ_1 in the interval $(p_{j-1}, p_j]$, z_2 some fixed value of τ_2 in the interval $(p_{j-1}, p_j]$, and m_1 a fixed value of $\mu_1^{p_j}$. We have

- for all $p_{j-1} < t < p_j$: $\Pr[\min_i c_i^t(x) + \mu_1^t \geq \text{Thresh}_1 + z_1] \leq \Pr[\min_i c_i^t(y) + \mu_1^t \geq \text{Thresh}_1 + z_1 - 1]$,
- for all $p_{j-1} < t < p_j$: $\Pr[\max_i c_i^t(x) + \mu_2^t \leq \text{Thresh}_2 + z_2] \leq \Pr[\max_i c_i^t(y) + \mu_2^t \leq \text{Thresh}_2 + z_2 + 1]$,
- for $t = p_j$: $\Pr[\min_i c_i^t(x) + m_1 < \text{Thresh}_1 + z_1] \leq \Pr[\min_i c_i^t(y) + m_1 - 2 < \text{Thresh}_1 + z_1 - 1]$.

Thus, the same outcome can be achieved by shifting τ_1 and τ_2 by at most 1, and $\mu_1^{p_j}$ by at most 2. By integrating in the same way as in the proof of Lemma B.1, since τ_1 and τ_2 are distributed according to $\text{Lap}(9/\epsilon)$ and μ_1^t is distributed according to $\text{Lap}(18/\epsilon)$, the distributions are $\epsilon/3$ -close.

Next, we analyze the probabilities that at time p_j , condition 34 was fulfilled, and neither 34 nor 22 was fulfilled at times $t \in (p_{j-1}, p_j)$, and the condition in line 22 was not fulfilled at time p_j (case B). Similar to before, let z_1 be some fixed value for τ_1 in the interval $(p_{j-1}, p_j]$, z_2 some fixed value of τ_2 in the interval $(p_{j-1}, p_j]$, and m_2 a fixed value of $\mu_2^{p_j}$. We have

- for all $p_{j-1} < t \leq p_j$: $\Pr[\min_i c_i^t(x) + \mu_1^t \geq \text{Thresh}_1 + z_1] \leq \Pr[\min_i c_i^t(y) + \mu_1^t \geq \text{Thresh}_1 + z_1 - 1]$,
- for all $p_{j-1} < t < p_j$: $\Pr[\max_i c_i^t(x) + \mu_2^t \leq \text{Thresh}_2 + z_2] \leq \Pr[\max_i c_i^t(y) + \mu_2^t \leq \text{Thresh}_2 + z_2 + 1]$,
- for $t = p_j$: $\Pr[\max_i c_i^t(x) + m_2 > \text{Thresh}_2 + z_2] \leq \Pr[\max_i c_i^t(y) + m_2 + 2 > \text{Thresh}_2 + z_2 + 1]$.

Thus, the same outcome can be achieved by shifting τ_1 and τ_2 by at most 1, and $\mu_2^{p_j}$ by at most 2. By integrating in the same way as in the proof of Lemma B.1, since τ_1 and τ_2 are distributed according to $\text{Lap}(9/\epsilon)$ and μ_2^t is distributed according to $\text{Lap}(18/\epsilon)$, the distributions are $\epsilon/3$ -close.

Conditioning on ending the j -th interval at p_j and case A resp. case B, we need to argue about the updating of the modes (lines 30 and 42, respectively). Since we add $\text{Lap}(3d/\epsilon)$ to $c_i^{p_j}(x)$ resp. $c_i^{p_j}(y)$ for all i , and $\|c^{p_j}(x) -$

$c^{p_j}(y)||_1 \leq d$, the probabilities of any output set are $\epsilon/3$ -close by the properties of the Laplace mechanism. By post-processing, the probabilities of updating $mode_i$ to M^{p_j} are $\epsilon/3$ -close on x and y . Together, the probabilities of getting $[p_0, p_1), \dots, [p_{j-1}, p_j)$ and M^0, \dots, M^{p_j} are $2\epsilon/3$ close on x and y . Since the rest of the mechanism depends only on p_j , M^{p_j} , and the input streams for times $t > p_j > t^*$, conditioning on p_j and M^{p_j} the probabilities are equal. We get that the probabilities of getting P and M are $2\epsilon/3$ -close on x and y . This shows that the partitioning mechanism together with the mode updates is $2\epsilon/3$ -differentially private. \square

C.2 Accuracy

Let K_t be the number of times up to time t that two consecutive input data rows differ, even if they differ just in one coordinate, i.e., the number of time such that $x^{t'} \neq x^{t'+1}$ for $t' \leq t$.

Lemma C.3. *With probability at least $1 - 3\beta$, Mechanism 6 has error at most $O(\text{err}(9K_t + 9, \beta) + \frac{d}{\epsilon} \log(9K_t + 9) + \frac{1}{\epsilon} \log(t/\beta))$ at all time steps t , where $\text{err}(\ell, \beta)$ is the error of the histogram mechanism H that holds with probability at least $1 - \beta$ for all length- ℓ prefixes of the input stream.*

Similarly to the accuracy proof in Section B, our proof of Lemma C.3 builds on first bounding the values of all random variables and the error of H such that all bounds hold simultaneously with probability $1 - 3\beta$. We call this event E and condition on it. Formally,

1. Let $t \in [1, T]$. With probability at least $1 - \beta_t/4$, any Y drawn from $\text{Lap}(b/\epsilon)$ for any b has an absolute value of at most $\frac{b \log(4/\beta_t)}{\epsilon}$. Thus, by a union bound, with probability at least $1 - \beta_t$, we have $|\mu_i^t| \leq \frac{18 \log(4/\beta_t)}{\epsilon}$ and $|\tau_i| \leq \frac{9 \log(4/\beta_t)}{\epsilon}$ for $i = 1$ and $i = 2$ at any fixed time t . Using a union bound over all time steps t and observing that $\sum_{t \in [1, T]} \beta_t = \sum_{t \in [1, T]} 6\beta' / (\pi^2 t^2) \leq \beta$, it follows that with probability at least $1 - \beta$ for $i = 1$ and $i = 2$ and for all time steps t that $|\mu_i^t| \leq \frac{18 \log(4/\beta_t)}{\epsilon}$.
2. With probability at least $1 - \beta_j$, any Y drawn from $\text{Lap}(3d/\epsilon)$ has an absolute value of at most $\frac{3d}{\epsilon} \log(1/\beta_j)$. Thus, by a union bound as above, with probability at least $1 - \beta$, we have $|\lambda_i| \leq \frac{3d}{\epsilon} \log(1/\beta_j)$ for all values of λ_i .
3. By the properties of H , with probability at least $1 - \beta$, the error of H after j inputs is at most $\text{err}(j, \beta)$ for all j .

Event E is the event that all three conditions hold, which happens with probability at least $1 - 3\beta$. The proof now consists of two main lemmata, Lemma C.5 and Lemma C.6. To show them we need:

Claim C.4. Conditioned on event E the following hold:

- 1) If at some time t the condition in line 22 is true for some $i \in [1, d]$, then the condition in line 29 is true for all $\ell \in [1, d]$ with $c_\ell = c_i$ and $mode_\ell = mode_i$, i.e., $mode_\ell$ is updated for all such ℓ .
- 2) If at some time t the condition in line 34 is true for some $i \in [1, d]$, then the condition in line 41 is true for all $\ell \in [1, d]$ with $c_\ell = c_i$ and $mode_\ell = mode_i$, i.e., $mode_\ell$ is updated for all such ℓ .

Proof of Claim C.4. If the condition in line 22 or line 34 is true, $t = p_j$ for some j . In the following, we use variable names to denote their value at time t when line 22 is reached. Further, if the condition in line 22 is true, then by the assumed bounds on the random variables in event E , there exists a c_i such that $c_i < \text{Thresh}_{i,1} + \frac{27 \log(4/\beta_t)}{\epsilon}$. Since $|\lambda_i| \leq \frac{3d}{\epsilon} \log(1/\beta_j)$, we have $c_i + \lambda_i \leq c_i + \frac{3d}{\epsilon} \log(1/\beta_j) < \text{Thresh}_{i,1} + \frac{27 \log(4/\beta_t)}{\epsilon} + \frac{3d}{\epsilon} \log(1/\beta_j) = \text{Thresh}_{i,1} + \alpha_t$, by definition of α_t . Thus, the condition in line 29 is true for i . Note that this is also case for all $\ell \in [1, d] \setminus i$ with $c_\ell = c_i$ and $mode_\ell = mode_i$, since then $\text{Thresh}_{i,1} = \text{Thresh}_{\ell,1}$. Now, to show that $mode_i$ actually changed, we need to show that $mode_i \neq -1$ at the beginning of the round. Assume $mode_i = -1$ at the beginning of the round. Then $c_i + \lambda_i < \text{Thresh}_{i,1} + \alpha_t = -t_{\text{diff}} - \alpha_t$, thus $c_i < -t_{\text{diff}} - \alpha_t + \frac{3d}{\epsilon} \log(1/\beta_j) < -t_{\text{diff}}$, which is a contradiction since $t_{\text{diff}} \geq c_i \geq -t_{\text{diff}}$ always. By the same argument, $mode_\ell$ is updated for all $\ell \in [1, d]$ with $c_\ell = c_i$ and $mode_\ell = mode_i$. This proves the first part of the claim.

Similarly, if the condition in line 34 is true, then by the assumed bounds on the random variables, there exists an i such that $c_i > \text{Thresh}_{i,2} - \frac{27 \log(4/\beta_t)}{\epsilon}$. Note that this is also case for all $\ell \in [1, d]$, $j \neq i$ with $c_j = c_i$. Since $|\lambda_i| \leq \frac{3d}{\epsilon} \log(1/\beta_j)$, we have $c_i + \lambda_i \geq c_i - \frac{3d}{\epsilon} \log(1/\beta_j) > \text{Thresh}_{i,2} - \frac{27 \log(4/\beta_t)}{\epsilon} - \frac{3d}{\epsilon} \log(1/\beta_j) = \text{Thresh}_{i,2} - \alpha_t$,

by definition of α_t . Thus, the condition in line 41 is true for i . Note that this is also case for all $\ell \in [1, d]$, $\ell \neq i$ with $c_\ell = c_i$ and $mode_\ell = mode_i$. Now, to show that $mode_i$ actually changed, we need to show that $mode_i \neq 1$ at the beginning of the round. Assume it was. Then $c_i + \lambda_i > \text{Thresh}_{i,2} - \alpha_t = t_{\text{diff}} + \alpha_t$, thus $c_i > t_{\text{diff}} + \alpha_t - \frac{3d}{\epsilon} \log(1/\beta_j) > t_{\text{diff}}$, which is a contradiction since $t_{\text{diff}} \geq c_i \geq -t_{\text{diff}}$ always. By the same argument, $mode_\ell$ is changed, for all $\ell \in [1, d]$ with $c_\ell = c_i$ and $mode_\ell = mode_i$. This proves the second part of the claim. \square

The next lemma shows an error bound at time t on the output of Mechanism 6 depending on the number of intervals (n_t) produced by the mechanism. This follows from (A) the fact that the histogram receives n_t inputs, and (B) from the fact that we can bound the additional error accumulated within an interval by $O(\frac{d}{\epsilon} \log(n_t) + \frac{1}{\epsilon} \log(t/\beta))$.

Lemma C.5. *Let n_t be the number of closed intervals at the end of time step t , i.e., if $t = p_j$ for some j , then $n_t = j$, and if $t \in (p_{j-1}, p_j)$, then $n_t = j - 1$. Conditioned on event E , the maximum additive error for all time steps $t' \leq t$ is $O(\text{err}(n_t, \beta) + \frac{d}{\epsilon} \log(n_t) + \frac{1}{\epsilon} \log(t/\beta))$.*

Proof of Lemma C.5. We differentiate two cases.

1. If $t = p_j$ for some $j \leq n_t$, the output is equal to H_{out} , in which case the error is at most $\text{err}(n_t, \beta)$, by the properties of H .
2. If $t \in (p_{j-1}, p_j)$, $j \leq n_t + 1$, let $t = p_{j-1} + t_{\text{diff}}$ and denote by out^t the output at time t . The error at time step t is given by $\max_i |\sum_{t'=1}^t x_i^{t'} - \text{out}^t| \leq \max_i (|\sum_{t'=1}^{p_{j-1}} x_i^{t'} - \text{out}^{p_{j-1}}| + |c_i - mode_i \cdot t_{\text{diff}}|)$, where c_i , $mode_i$ and t_{diff} correspond to the values of those variables in the mechanism at time t . $|\sum_{t'=1}^{p_{j-1}} x_i^{t'} - \text{out}^{p_{j-1}}| \leq \text{err}(n_t, \beta)$ by the properties of H . Thus, we bound $|c_i - mode_i \cdot t_{\text{diff}}|$. If either condition 34 or 22 would have been true, then $t = p_j$ for some j , a contradiction. Thus, both conditions were false. Since we conditioned on the absolute values of μ_1^t, μ_2^t being bounded by $\frac{18 \log(4/\beta_t)}{\epsilon}$ and the absolute values of τ_1, τ_2 being bounded by $\frac{9 \log(4/\beta_t)}{\epsilon}$, we have $c_i - mode_i \cdot t_{\text{diff}} + 2\alpha_t \geq -\frac{27 \log(4/\beta_t)}{\epsilon}$ for all i (because 22 is false) and $c_i - mode_i \cdot t_{\text{diff}} - 2\alpha_t \leq \frac{27 \log(4/\beta_t)}{\epsilon}$ for all i (because 34 is false). We have $|c_i - mode_i \cdot t_{\text{diff}}| \leq \frac{27 \log(4/\beta_t)}{\epsilon} + 2\alpha_t = O(\frac{d}{\epsilon} \log(1/\beta_j) + \frac{1}{\epsilon} \log(1/\beta_t)) = O(\frac{d}{\epsilon} \log(n_t/\beta) + \frac{1}{\epsilon} \log(t/\beta))$. \square

Lemma C.6. *Conditioned on event E , at any time step t , no more than $\min(t, 9K_t + 9)$ intervals were closed at the end of time step t .*

To prove Lemma C.6 we partition the stream of input rows into *episodes* such that (1) an episode starts at time $t = 1$ and also at time $t' \in [T]$ if row $x^{t'-1}$ and row $x^{t'}$ differ, and (2) an episode ends at the end of the stream and also at time $t' \in [T]$ if row $x^{t'}$ and row $x^{t'+1}$ differ. Our proof idea is to show that in no episode more than 9 intervals are closed. As there are at most $K_t + 1$ episodes by the definition of episodes, Lemma C.6 will follow. For episodes in which no interval is closed, nothing has to be shown. Thus, we study in the following episodes in which at least one interval is closed. We first show the following claims, which basically show that if the mechanism receives the same row x^t for a “long enough” time period, then eventually it will set the mode vector equal to x^t .

Let c_i^t (resp. $\text{Thresh}_{i,1}^t$ resp. $\text{Thresh}_{i,2}^t$) denote the value of c_i (resp. $\text{Thresh}_{i,1}$ resp. $\text{Thresh}_{i,2}$) after the initial processing of time step t , i.e., in line 22.

Claim C.7. Let I be an episode in which at least one interval is closed and let t^* be the first time step at which an interval is closed in I . Conditioned on event E , if at any time step $t > t^*$ where $t \in I$ row x^t equals the vector $mode$ then no mode is updated in time step t .

Proof. Consider a time step $t > t^*$ in I and let t' with $t^* \leq t' < t$ be the last time before t that a mode was changed. Then c_i is reset to 0 at time t' , for all $i \in [d]$. By definition of t' , $mode_i$ did not change since t' , and by definition of I , x did not change since t' . Thus, $c^t = (t - t') \cdot x^t = t_{\text{diff}} \cdot x^t = t_{\text{diff}} \cdot mode$.

Thus for each $i \in [d]$, $c_i^t - \text{Thresh}_{i,1}^t = 2\alpha_t$ and $c_i^t - \text{Thresh}_{i,2}^t = -2\alpha_t$ for all $i \in [d]$. It follows that $\min_i (c_i^t - \text{Thresh}_{i,1}^t) = 2\alpha_t$ and $\min_i (c_i^t - \text{Thresh}_{i,2}^t) = -2\alpha_t$. But since we condition on E , $|\mu_1^t| \leq \frac{18 \log(4/\beta_t)}{\epsilon}$ and $|\tau_1| \leq \frac{9 \log(4/\beta_t)}{\epsilon}$, it follows that $2\alpha_t = \frac{54}{\epsilon} \log(4/\beta_t) + \frac{6d}{\epsilon} \log(1/\beta_j) > \tau_1 - \mu_1^t$ and that $-2\alpha_t = -\frac{54}{\epsilon} \log(4/\beta_t) -$

$\frac{6d}{\epsilon} \log(1/\beta_j) < \tau_2 - \mu_2^t$. Thus, the conditions on lines 22 and 34 cannot hold and the interval does not close at time t . \square

Claim C.8. Let I be an episode in which at least one interval is closed and let t^* be the first time step at which an interval is closed in I . Conditioned on event E , at any time step $t > t^*$ where $t \in I$ and for any $i \in [d]$ if $mode_i < x_i^t$ at the start of t then $mode_i$ will not decrease in time step t and, symmetrically, if $mode_i > x_i^t$ at the start of t then $mode_i$ will not increase in time step t .

Proof. Consider a coordinate $i \in [d]$ and a time step $t > t^*$ in I and let t' with $t^* \leq t' < t$ be the last time before t that a mode was changed. Then c_i is reset to 0 at time t' , for all $i \in [d]$. By definition of t' $mode_i$ did not change since t' , and by definition of I , x_i did not change since t' . Thus, $c_i^t = (t - t') \cdot x_i^t = t_{\text{diff}} \cdot x_i^t$.

If $x_i^t > mode_i$, then $c_i^t \geq t_{\text{diff}} \cdot mode_i$. Due to the conditioning it follows that $c_i^t - \text{Thresh}_{i,1}^t = c_i^t - t_{\text{diff}} \cdot mode_i + 2\alpha_t \geq 2\alpha_t = \frac{27}{\epsilon} \log(4/\beta_t) + \frac{3d}{\epsilon} \log(1/\beta_j) + \alpha_t > \alpha_t - \lambda_i$, and, thus, the condition in Line 29 does not hold in time step t and $mode_i$ will not decrease.

If $x_i^t < mode_i$, then $c_i^t \leq t_{\text{diff}} \cdot mode_i$. Due to the conditioning it follows that $c_i^t - \text{Thresh}_{i,2}^t = c_i^t - t_{\text{diff}} \cdot mode_i - 2\alpha_t \leq -2\alpha_t = -\frac{27}{\epsilon} \log(4/\beta_t) - \frac{3d}{\epsilon} \log(1/\beta_j) - \alpha_t < -\alpha_t - \lambda_i$, and, thus, the condition in Line 41 does not hold in time step t and $mode_i$ will not increase. \square

The proof of Lemma C.6 now consists of a careful case analysis using Claim C.7 and Claim C.8 to show that within any episode, at most 9 intervals will be closed.

Proof of Lemma C.6. The claim that there are at most t closed intervals follows trivially as at most one interval is closed in a single time step.

We proceed to show that there are at most $9K_t + 9$ closed intervals up to time step t . By Claim C.4 whenever an interval ends, at least one mode has to change. Thus, we will study in the following how many time steps exist that contain a mode update. By the definition of K_t there are exactly $K_t + 1$ many episodes up to time step t and, thus, it suffices to show that for any episode I , there are at most 9 time steps in I where a mode is updated.

If no interval is closed in I , i.e., no mode is ever updated, the claim holds trivially for I . Thus in the following assume that there is at least one time step where a mode is updated and let t^* be the first such time step. As a shorthand we use m_i to denote the value of $mode_i$ at the end of time step t^* . Note that $c_i^{t^*} = 0$ and $x_i^t = x_i^{t^*}$ for all $i \in [d]$ and all $t \in I$. Thus for all subsequent time steps $t > t^*$ in I and for all coordinates i, j with $x_i^t = x_j^t$ it holds that $c_i^t = c_j^t$. Thus, by Claim C.4, every time t an interval is closed, there exists an $x^* \in \{-1, 0, 1\}$ and $m^* \in \{-1, 0, 1\}$ such that $mode_i$ is updated for all i with $x_i^t = x^*$ and $m_i = m^*$. As there are only 3 possible values that a variable $x_i^t = x_i^{t^*}$ can assume, it suffices to study for each value $x^* \in \{-1, 0, 1\}$ how often a coordinate i with $x_i^{t^*} = x^*$ updates its mode within I . We analyze three cases:

First we consider all coordinates i with $x_i^{t^*} = -1$ and partition them into 3 subgroups depending on their m_i value. For $m_i = -1$, Claim C.7 shows that there are no time steps with mode updates as the value of the mode and x_i are equal. For $m_i = 0$, Claim C.8 shows that there is at most one time step with mode updates, which decreases the mode to -1. For $m_i = 1$, Claim C.8 shows that there are at most two time steps with mode updates, each decreasing the mode by 1. Thus there are at most 3 time steps with mode updates for all coordinates i with $x_i^{t^*} = -1$.

Next consider all coordinates i with $x_i^{t^*} = 0$ and partition them into 3 subgroups depending on m_i . For $m_i = 0$, there are no time steps with mode updates as the value of the mode and x_i are equal. For $m_i = 1$, there is at most one time step with mode updates, which decreases the mode to 0. For $m_i = -1$, there is at most one time steps with mode updates, which increases the mode to 0. Thus there are at most 2 time steps with mode updates for all coordinates i with $x_i^{t^*} = 0$.

Finally we consider all coordinates i with $x_i^{t^*} = 1$ and partition them into 3 subgroups depending on m_i . For $m_i = 1$, Claim C.7 shows that there are no time steps with mode updates as the value of the mode and x_i are equal. For $m_i = 0$, there is at most one time step with mode updates, which increases the mode to 1. For $m_i = -1$, there are at most two time steps with mode updates, each increasing the mode by 1. Thus there are at most 3 time steps with mode updates for all coordinates i with $x_i^{t^*} = 1$.

Thus, combined with the update at time step t^* a mode update happens in at most 9 time steps in episode I . This concludes the proof. \square

Lemma C.3 now follows by Lemma C.5 and Lemma C.6.

D AN $\Omega(d \cdot \log T)$ LOWER BOUND FOR INDEPENDENTLY DIFFERENTIALLY PRIVATE d -DIMENSIONAL BINARY COUNTING

So far, all upper bounds for HISTOGRAM and even for MAXSUM which were not polynomial in T relied on running d independent binary counting mechanisms in parallel, and all achieved an error $\Omega(d \log T)$ for ϵ -differential privacy. In this section we prove that using this strategy one cannot do better. For this, we formally define the following alternative version of differential privacy and neighboring streams of elements from $\{0, 1\}^d$:

Definition D.1 (Independent differential privacy). Let x be a stream of T elements from $\{0, 1\}^d$. We say x and y are *independently neighboring* if and only if for every $i \in [1, d]$ there exists a time step t_i such that $x_i^t = y_i^t$ for all $t \neq t_i$. A mechanism A is *independently ϵ -differentially private* if it fulfills Definition 3.2 for independently neighboring x and y .

Note that this is a superset of the earlier definition of neighboring streams, i.e., all x and y which are neighboring are also independently neighboring, but not vice-versa. Thus, independent differential privacy is a stronger property than differential privacy.

We show the lower bound using a *packing argument* (Hardt and Talwar, 2010), which relies on the *group privacy* property of differential privacy summarized in Fact 6. We say x and y are k -neighboring if there exist $x = X_1, X_2, \dots, X_k = y$ such that X_i and X_{i+1} are neighboring for all $1 \leq i < k$. In the same way, we say x and y are *independently k -neighboring* if there exist $x = X_1, X_2, \dots, X_k = y$ such that X_i and X_{i+1} are independently neighboring for all $1 \leq i < k$.

Fact 6. Let A be an ϵ -(independently) differentially private mechanism and x and y be (independently) k -neighboring. Then for all $S \in \text{range}(A)$

$$P(A(x) \in S) \leq e^{k\epsilon} P(A(y) \in S)$$

Note that computing all d column sums by running independent binary counting mechanisms fulfills independent ϵ -differential privacy. Next we show that $\Omega(d \log T)$ noise is necessary for computing the noisy column sums in every time step while preserving independent ϵ -differential privacy.

Theorem 5. Assume $d \leq \sqrt{T}$. Then there is a $T' = T'(\epsilon)$ such that any independently ϵ -differentially private mechanism for computing all d column sums cannot be (α, β) -accurate for constant β and $\alpha \leq \frac{d \ln T}{16\epsilon}$ for streams of length $T \geq T'$.

Proof. Let $b = \frac{d \ln T}{8\epsilon}$. We assume without loss of generality that T is a multiple of b , otherwise we pad the stream with zero vectors. We start by dividing $[1, T]$ into T/b blocks of length b , i.e. $B_1 = [1, b]$, $B_2 = [b + 1, 2b]$, \dots , $B_{T/b} = [T - b + 1, T]$. Now, for any vector $v \in [T/b]^d$ define the following stream $x(v)$ and output data set $S(v)$:

- $x(v)_i^t = 1$ if and only if $t \in [(v_i - 1) \cdot b + 1, v_i \cdot b]$, else $x(v)_i^t = 0$. That is, for every coordinate d there is exactly one block which consists of only ones, and that block is the one specified by the i th coordinate in v .
- Denote s_i^t the estimate for the i th column sum output by the mechanism in time t . $S(v)$ includes all outputs such that $s_i^t < b/2$ for all $t \leq (v_i - 1) \cdot b$ and $s_i^t > b/2$ for all $t \geq v_i \cdot b$.

Now let $\beta = 1/3$ and assume there is a mechanism A which is (α, β) accurate for $\alpha = \frac{d \ln T}{16\epsilon} = b/2$. Then $P(A(x(v)) \in S(v)) \geq \frac{2}{3}$. Further, we have that for any $v, v' \in [T/b]^d$, $x(v)$ and $x(v')$ are independently $2b$ -neighboring: for every coordinate i , $x(v')_i$ and $x(v)_i$ differ in at most $2b$ timesteps. Therefore, by Fact 6,

$P(A(x(v)) \in S(v')) \geq \frac{2}{3}e^{-2b\epsilon}$. Additionally, all $S(v)$ are disjoint for distinct v . Thus

$$\begin{aligned} 1 &\geq \sum_{v \in [T/b]^d} \frac{2}{3}e^{-2b\epsilon} \\ &= \frac{2}{3} \left(\frac{8T\epsilon}{d \ln T} \right)^d e^{-\frac{d \ln T}{4}} \\ &\geq \frac{2}{3} \left(\frac{8\sqrt{T}\epsilon}{\ln T} \right)^d (1/\sqrt[4]{T})^d \\ &= \frac{2}{3} \left(\frac{8\sqrt[4]{T}\epsilon}{\ln T} \right)^d > 1 \end{aligned}$$

for large enough T , which is a contradiction. \square

E USAGE OF CONCURRENT COMPOSITION

Our ϵ -dp result could alternatively be shown as follows: One could use the result of [Qiu and Yi \(2022\)](#) to argue adaptive parallel composition for the partitioning mechanism, then Fact 4 by [Denisov et al. \(2022\)](#) to argue that the ϵ -dp partitioning mechanism in the continual release model is also ϵ -dp in the adaptive continual release model, and then use the result of [Vadhan and Wang \(2021\)](#) to concurrently compose the adaptive partitioning and adaptive histogram mechanisms. However, this would not reduce the technical complexity of the proof, and also not be self-contained.

Moreover, this proof strategy does not work for (ϵ, δ) -dp at all. Adaptive parallel composition in the (ϵ, δ) -dp setting is an open problem. [Guerra-Balboa et al. \(2024\)](#) give an adaptive parallel composition theorem for (ϵ, δ) -dp, but their result assumes that the partition of the dataset is given beforehand, while we require that the partition of the dataset is also performed adaptively. Further, it is not enough to show that the partitioning mechanism is differentially private, we would need to show that it is *adaptively* private since a general transformation as in the ϵ -dp case does not exist here, and is only known for specific mechanisms.

F EXTENSION TO (ϵ, δ) -DIFFERENTIAL PRIVACY

We will use noise drawn from the Normal distribution for our mechanism. The mechanism constructed using noise drawn from the Normal distribution is known as the Gaussian mechanism, which satisfies (ϵ, δ) -dp.

Definition F.1 (Normal Distribution). The *normal distribution* centered at 0 with variance σ^2 is the distribution with the probability density function

$$f_{N(0, \sigma^2)}(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right)$$

We use $X \sim N(0, \sigma^2)$ or sometimes just $N(0, \sigma^2)$ to denote a random variable X distributed according to $f_{N(0, \sigma^2)}$.

Fact 7 (Theorem A.1 in [Dwork and Roth \(2014\)](#): Gaussian mechanism). Let f be any function $f : \mathcal{X} \rightarrow \mathbb{R}^m$ with L_2 -sensitivity Δ_2 . Let $\epsilon \in (0, 1)$, $c^2 > 2 \ln(1.25/\delta)$, and $\sigma \geq c\Delta_2(f)/\epsilon$. Let $Y_i \sim N(0, \sigma^2)$ for $i \in [m]$. Then the mechanism defined as:

$$A(x) = f(x) + (Y_1, \dots, Y_m)$$

satisfies (ϵ, δ) -differential privacy.

We use the following continual histogram mechanism H introduced by [Fichtenberger et al. \(2023\)](#), which achieves an error of $O(\epsilon^{-1} \log(1/\delta) \log t \sqrt{d \ln(dt)})$ at time step t . Since their mechanism fulfills the conditions of Theorem 2.1 of [Denisov et al. \(2022\)](#), the same privacy guarantees hold for their mechanism in the adaptive continual release model.

Fact 8 ((ϵ, δ) -differentially private continual histogram against an adaptive adversary). There is an (ϵ, δ) -differentially private mechanism in the adaptive continual release model for continual histogram that with probability $\geq 1 - \beta$, has error bounded by $O(\epsilon^{-1} \log(1/\delta) \log t \sqrt{d \ln(dt/\beta)})$ at time t .

F.1 Histogram Queries

We make the following changes to the mechanism to obtain an (ϵ, δ) -dp mechanism for histogram queries.

1. Initialize an $(\epsilon/3, \delta/(2e^{2\epsilon/3}))$ -adaptively dp continual histogram mechanism H .
2. Sample $\gamma_k^j \sim N(0, 18k \ln(4e^{2\epsilon/3}/\delta)/\epsilon^2)$.
3. Set α_γ^j to $6\epsilon^{-1}\sqrt{m \ln(12e^{2\epsilon/3}m/(\delta\beta_j))}$.

Privacy. We detail the changes to the privacy proof from the ϵ -dp case. As in the ϵ -dp case, we need to now show that

$$\Pr[\mathcal{A}(x) \in S] \leq e^\epsilon \cdot \Pr[\mathcal{A}(y) \in S] + \delta$$

Since H is $(\epsilon/3, \delta/(2e^{2\epsilon/3}))$ -adaptively differentially private, we get that

$$\Pr(V_{H, Adv(x,y)}^{(x)} \in S) \leq e^{\epsilon/3} \Pr(V_{H, Adv(x,y)}^{(y)} \in S) + \delta/(2e^{2\epsilon/3})$$

and

$$\Pr(V_{H, Adv(x,y)}^{(y)} \in S) \leq e^{\epsilon/3} \Pr(V_{H, Adv(x,y)}^{(x)} \in S) + \delta/(2e^{2\epsilon/3}).$$

Thus all we would need to show would be

$$\Pr(V_{H, Adv(x,y)}^{(x)} \in S) \leq e^{2\epsilon/3} \Pr(V_{H, Adv(y,x)}^{(x)} \in S) + \delta/2, \quad (3)$$

since then

$$\begin{aligned} \Pr(\mathcal{A}(x) \in S) &= \Pr(V_{H, Adv(x,y)}^{(x)} \in S) \\ &\leq e^{2\epsilon/3} \Pr(V_{H, Adv(y,x)}^{(x)} \in S) + \delta/2 \\ &\leq e^\epsilon \Pr(V_{H, Adv(y,x)}^{(y)} \in S) + \delta \\ &= e^\epsilon \Pr(\mathcal{A}(y) \in S) + \delta \end{aligned} \quad (4)$$

The partitioning is still $e^{\epsilon/3}$ -close by the same arguments since we use the same random variables as in the ϵ -dp case. For the thresholds, note that conditioned on all previous outputs of H and p_j being equal, $q_k(s^{p_j}(x))$ and $q_k(s^{p_j}(y))$ can differ by at most 1 for each $k \in [m]$. Thus the L_2 difference between the two vectors is at most \sqrt{m} . By Lemma 7 for the Gaussian mechanism, adding $N(0, 18k \ln(4e^{2\epsilon/3}/\delta)/\epsilon^2)$ noise to every $q_k(s^{p_j}(y))$ ensures that the distributions of $q_k(s^{p_j}(x)) + \gamma_k^j$ and $q_k(s^{p_j}(y)) + \gamma_k^j$ are $(e^{\epsilon/3}, \delta/2e^{2\epsilon/3})$ -close for all $k \in [m]$. Since the condition in line 31 only depends on those, this implies that the probabilities of executing line 31 on any subset of $[m]$ on $\text{run}(x)$ and $\text{run}(y)$ are $(e^{\epsilon/3}, \delta/2e^{\epsilon/3})$ -close, as required.

Accuracy. We have that Lemma B.2 holds with $\alpha_\mu^t, \alpha_\tau^j$ as earlier, $\alpha_\gamma^j = 6\epsilon^{-1}\sqrt{m \ln(12e^{2\epsilon/3}m/(\delta\beta_j))}$ and $\alpha_H^j = O(\epsilon^{-1} \log(1/\delta) \log j \sqrt{d \ln(dj/\beta)})$. Thus, by Lemma B.7, the mechanism has error at most

$$\alpha^{(t)} = O\left(\epsilon^{-1} \log(1/\delta) \cdot \left(\sqrt{d} \log^{3/2}(dmq^{*t}/\beta) + \sqrt{m} \log(mq^{*t}/\beta) + \log t\right)\right)$$

at all time steps t with probability at least $1 - \beta$ as required.

F.2 Histogram Parameterized in the Number of Fluctuations

We make the following changes to the mechanism to obtain an (ϵ, δ) -dp mechanism for HISTOGRAM parameterized in the number of fluctuations.

1. Initialize an $(\epsilon/3, \delta/2)$ -adaptively dp continual histogram mechanism H .
2. Sample $\gamma_i^j \sim N(0, 18d \ln(4e^{2\epsilon/3}/\delta)/\epsilon^2)$.
3. Replace $\frac{3d}{\epsilon} \log(1/\beta_j)$ with $6\epsilon^{-1}\sqrt{d \ln(4e^{2\epsilon/3}/\delta\beta_j)}$.

Privacy. We detail the changes to the privacy proof from the ϵ -dp case.

We will show that the computation of p_0, p_1, \dots along with the sequence of *mode* value updates is $(2\epsilon/3, \delta/2)$ -dp. Then the result follows by basic composition with the histogram mechanism.

The partitioning is still $e^{\epsilon/3}$ -close by the same arguments since we use the same random variables as in the ϵ -dp case. For the modes, note that conditioning on ending the j -th interval at p_j , $\|c^{p_j}(x) - c^{p_j}(y)\|_2 \leq \sqrt{d}$. Thus, by Lemma 7, adding $N(0, 18d \ln(4e^{2\epsilon/3}/\delta)/\epsilon^2)$ noise to each $c_i^{p_j}$ ensures that the distributions on x and y are $(e^{\epsilon/3}, \delta/2e^{\epsilon/3})$ -close. Thus both the partitioning and mode updates together are $(2\epsilon/3, \delta/2)$ -dp as required.

Accuracy. Using the histogram mechanism from Fact 8, and replacing $\frac{d}{\epsilon} \log(1/\beta_j)$ in the proof with $6\epsilon^{-1} \sqrt{d \ln(4e^{2\epsilon/3}/\delta\beta_j)}$, we get that the mechanism has error at most

$$O\left(\epsilon^{-1} \log(1/\delta) \cdot \left(\sqrt{d} \log^{3/2}(dK/\beta) + \log t\right)\right)$$

at all time steps t with probability at least $1 - \beta$.

G THE SPARSE VECTOR TECHNIQUE

The sparse vector technique is based on an algorithm by Dwork et al. (2009) and was described more fully by Dwork and Roth (2014). The version described in Algorithm 7 is by Lyu et al. (2017) for $c = 1$ (the main difference is that it allows different thresholds for every query).

Mechanism 7 AboveThreshold

```

1: Input: Data Set  $D$ , Sensitivity bound  $\Delta$ , thresholds  $\text{Thresh}_1, \text{Thresh}_2, \dots$ , and queries  $q_1, q_2, \dots$  which are
   have sensitivity at most  $\Delta$ 
2:  $\tau = \text{Lap}(2\Delta/\epsilon)$ 
3: for  $i = 1, \dots$ , do
4:    $\mu_i = \text{Lap}(4\Delta/\epsilon)$ 
5:   if  $q_i(D) + \mu_i > \text{Thresh}_i + \tau$  then
6:     output  $a_i = \text{YES}$ 
7:     Abort
8:   else
9:     output  $a_i = \text{NO}$ 

```

Lemma G.1 (Dwork and Roth, 2014; Lyu et al., 2017). *Algorithm 7 is ϵ -differentially private.*

Lemma G.2 (Dwork and Roth, 2014; Lyu et al., 2017). *Algorithm 7 fulfills the following accuracy guarantees for $\alpha = \frac{8(\ln k + \ln(2/\beta))}{\epsilon}$: For any sequence q_1, \dots, q_k of queries it holds with probability at least $1 - \beta$,*

1. *for i such that $a_i = \text{YES}$ we have*

$$q_i(D) \geq \text{Thresh}_i - \alpha,$$

2. *for all i such that $a_i = \text{NO}$ we have*

$$q_i(D) \leq \text{Thresh}_i + \alpha.$$