Local Differential Privacy for Evolving Data

Matthew Joseph Computer and Information Science University of Pennsylvania majos@cis.upenn.edu

Jonathan Ullman Computer and Information Sciences Northeastern University jullman@ccs.neu.edu Aaron Roth Computer and Information Science University of Pennsylvania aaroth@cis.upenn.edu

Bo Waggoner Computer and Information Science University of Pennsylvania bowaggoner@gmail.com

Abstract

There are now several large scale deployments of differential privacy used to collect statistical information about users. However, these deployments periodically recollect the data and recompute the statistics using algorithms designed for a single use. As a result, these systems do not provide meaningful privacy guarantees over long time scales. Moreover, existing techniques to mitigate this effect do not apply in the "local model" of differential privacy that these systems use.

In this paper, we introduce a new technique for local differential privacy that makes it possible to maintain up-to-date statistics over time, with privacy guarantees that degrade only in the number of changes in the underlying distribution rather than the number of collection periods. We use our technique for tracking a changing statistic in the setting where users are partitioned into an unknown collection of groups, and at every time period each user draws a single bit from a common (but changing) group-specific distribution. We also provide an application to frequency and heavy-hitter estimation.

1 Introduction

After over a decade of research, differential privacy [12] is moving from theory to practice, with notable deployments by Google [15, 6], Apple [2], Microsoft [10], and the U.S. Census Bureau [1]. These deployments have revealed gaps between existing theory and the needs of practitioners. For example, the bulk of the differential privacy literature has focused on the *central model*, in which user data is collected by a trusted aggregator who performs and publishes the results of a differentially private computation [11]. However, Google, Apple, and Microsoft have instead chosen to operate in the *local model* [15, 6, 2, 10], where users individually randomize their data on their own devices and send it to a potentially untrusted aggregator for analysis [18]. In addition, the academic literature has largely focused on algorithms for performing one-time computations, like estimating many statistical quantities [7, 22, 16] or training a classifier [18, 9, 4]. Industrial applications, however have focused on tracking statistics about a user population, like the set of most frequently used emojis or words [2]. These statistics evolve over time and so must be re-computed periodically.

Together, the two problems of periodically recomputing a population statistic and operating in the local model pose a challenge. Naïvely repeating a differentially private computation causes the privacy loss to degrade as the square root of the number of recomputations, quickly leading to enormous values of ϵ . This naïve strategy is what is used in practice [15, 6, 2]. As a result, Tang et al. [23] discovered that the privacy parameters guaranteed by Apple's implementation of differentially private data collection

32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada.

can become unreasonably large even in relatively short time periods.¹ Published research on Google and Microsoft's deployments suggests that they encounter similar issues [15, 6, 10].

On inspection the naïve strategy of regular statistical updates seems wasteful as aggregate population statistics don't change very frequently—we expect that the most frequently visited website today will typically be the same as it was yesterday. However, population statistics do eventually change, and if we only recompute them infrequently, then we can be too slow to notice these changes.

The *central model* of differential privacy allows for an elegant solution to this problem. For large classes of statistics, we can use the *sparse vector technique* [13, 22, 16, 11] to repeatedly perform computations on a dataset such that the error required for a fixed privacy level grows not with the number of recomputations, but with the number of times the computation's outcome *changes significantly*. For statistics that are relatively stable over time, this technique dramatically reduces the overall error. Unfortunately, the sparse vector technique provably has no local analogue [18, 24].

In this paper we present a technique that makes it possible to repeatedly recompute a statistic with error that decays with the number of times the statistic changes significantly, rather than the number of times we recompute the current value of the statistic, all while satisfying local differential privacy. This technique allows for tracking of evolving local data in a way that makes it possible to quickly detect changes, at modest cost, so long as those changes are relatively infrequent. Our approach guarantees privacy under any conditions, and obtains good accuracy by leveraging three assumptions: (1) each user's data comes from one of m evolving distributions; (2), these distributions change relatively infrequently; and (3) users collect a certain amount of data during each reporting period, which we term an *epoch*. By varying the lengths of the epochs (for example, collecting reports hourly, daily, or weekly), we can trade off more frequent reports versus improved privacy and accuracy.

1.1 Our Results and Techniques

Although our techniques are rather general, we first focus our attention on the problem of privately estimating the average of bits, with one bit held by each user. This simple problem is widely applicable because most algorithms in the local model have the following structure: on each individual's device, data records are translated into a short bit vector using sketching or hashing techniques. The bits in this vector are perturbed to ensure privacy using a technique called *randomized response*, and the perturbed vector is then sent to a server for analysis. The server collects the perturbed vectors, averages them, and produces a data structure encoding some interesting statistical information about the users as a whole. Thus many algorithms (for example, those based on *statistical queries*) can be implemented using just the simple primitive of estimating the average of bits.

We analyze our algorithm in the following probabilistic model (see Section 3 for a formal description). The population of n users has an unknown partition into subgroups, each of which has size at least L, time proceeds in rounds, and in each round each user samples a private bit independently from their subgroup-specific distribution. The private data for each user consists of the vector of bits sampled across rounds, and our goal is to track the total population mean over time. We require that the estimate be private, and ask for the strong (and widely known) notion of *local differential privacy*—for every user, no matter how other users or the server behave, the distribution of the messages sent by that user should not depend significantly on that user's private data.

To circumvent the limits of local differential privacy, we consider a slightly relaxed estimation guarantee. Specifically, we batch the rounds into T epochs, each consisting of ℓ rounds, and aim in each epoch t to estimate p^t , the population-wide mean across the subgroups and rounds of epoch t. Thus, any sufficiently large changes in this mean will be identified after the current epoch completes, which we think of as introducing a small "delay".

Our main result is an algorithm that takes data generated according to our model, guarantees a fixed level of local privacy ε that grows (up to a certain point) with the number of distributional changes rather than the number of epochs, and guarantees that the estimates released at the end of each epoch are accurate up to error that scales sublinearly in $1/\ell$ and only polylogarithmically with the total number of epochs T. Our method improves over the naïve solution of simply recomputing the statistic every epoch – which would lead to either privacy parameter or error that scales linearly with the

¹Although the value of ϵ that Apple guarantees over the course of say, a week, is not meaningful on its own, Apple does take additional heuristic steps (as does Google) that make it difficult to combine user data from multiple data collections [2, 15, 6]. Thus, they may still provide a strong, if heuristic, privacy guarantee.

number of epochs—and offers a quantifiable way to reason about the interaction of collection times, reporting frequency, and accuracy. We note that one can alternatively phrase our algorithm so as to have a fixed error guarantee, and a *privacy cost* that scales dynamically with the number of times the distribution changes².

Theorem 1.1 (Protocol for Bernoulli Means, Informal Version of Theorem 4.3). In the above model, there is an ε -differentially private local protocol that achieves the following guarantee: with probability at least $1-\delta$, while the total number of elapsed epochs t where some subgroup distribution

has changed is fewer than $\varepsilon \cdot \min\left(\frac{L}{\sqrt{n\ln(mT/\delta)}}, \ln(T)\sqrt{\frac{n}{\ell}}\right)$, the protocol outputs estimates \tilde{p}^t where $|\tilde{n}^t - n^t| = O\left(\ln(T)\sqrt{\frac{\ln(nT/\delta)}{2}}\right)$

$$|\tilde{p}^t - p^t| = O\left(\ln(T)\sqrt{\frac{\ln(nT/\delta)}{\ell}}\right)$$

where L is the smallest subgroup size, n is the number of users, ℓ is the chosen epoch length, and T is the resulting number of epochs.

To interpret the theorem, consider the setting where there is only one subgroup and L = n. Then to achieve error α we need, ignoring log factors, $\ell \ge 1/\alpha^2$ and that fewer than $\varepsilon \alpha \sqrt{n}$ changes have occured. We emphasize that our algorithm satisfies ε -differential privacy for all inputs without a distributional assumption-only accuracy relies on distributional assumptions.

Finally, we demonstrate the versatility of our method as a basic building block in the design of locally differentially private algorithms for evolving data by applying it to the well-known heavy *hitters* problem. We do so by implementing a protocol due to [3] on top of our simple primitive. This adapted protocol enables us to efficiently track the evolution of histograms rather than single bits. Given a setting in which each user in each round independently draws an object from a discrete distribution over a dictionary of d elements, we demonstrate how to maintain a *frequency oracle* (a computationally efficient representation of a histogram) for that dictionary with accuracy guarantees that degrade with the number of times the distribution over the dictionary changes, and only polylogarithmically with the number of rounds. We summarize this result below.

Theorem 1.2 (Protocol for Heavy-Hitters, Informal Version of Theorem 5.2). In the above model, there is an ε -differentially private local protocol that achieves the following guarantee: with probability at least $1 - \delta$, while the total number of elapsed epochs t where some subgroup distribution

has changed is fewer than $\varepsilon \cdot \min\left(\frac{L}{\sqrt{n \ln(mT/\delta)}}, \ln(T)\sqrt{\frac{n \ln(nT/\delta)}{\ell}}\right)$ the protocol outputs estimate oracles \hat{f}^t such that for all $v \in [d]$

$$|\hat{f}^t(v) - \mathcal{P}^t(v)| = O\left(\ln(T)\sqrt{\frac{\ln(nT/\delta)}{\ell}} + \sqrt{\frac{\ln(dnT/\delta)}{n}}\right)$$

where n is the number of users, L is the smallest subgroup size, \mathcal{P}^t is the mean distribution over dictionary elements in epoch t, d is the number of dictionary elements, ℓ is the chosen epoch length, and T is the resulting number of epochs.

1.2 Related Work

The problem of privacy loss for persistent local statistics has been recognized since at least the original work of Erlingsson et al. [15] on RAPPOR (the first large-scale deployment of differential privacy in the local model). Erlingsson et al. [15] offers a heuristic memoization technique that impedes a certain straightforward attack but does not prevent the differential privacy loss from accumulating linearly in the number of times the protocol is run. Ding et al. [10] give a formal analysis of a similar memoization technique, but the resulting guarantee is not differential privacy—instead it is a privacy guarantee that depends on the behavior of other users, and may offer no protection to users with idiosyncratic device usage. In contrast, we give a worst-case differential privacy guarantee.

Our goal of maintaining a persistent statistical estimate is similar in spirit to the model of *privacy* under continual observation Dwork et al. [14]. The canonical problem for differential privacy under

²We can achieve a dynamic, data-dependent privacy guarantee using the notion of *ex-post* differential privacy [19], for example by using a so-called privacy odometer [21].

continual observation is to maintain a running count of a stream of bits. However, the problem we study is quite different. In the continual observation model, new users are arriving, while existing users' data does not change. In our model each user receives new information in each round. (Also, we work in the local model, which has not been the focus of the work on continual observation.)

The local model was originally introduced by Kasiviswanathan et al. [18], and the canonical algorithmic task performed in this model has become frequency estimation (and heavy hitters estimation). This problem has been studied in a series of theoretical [17, 3, 5, 8, 2] and practical works [15, 6, 2].

2 Local Differential Privacy

We require that our algorithms satisfy *local differential privacy*. Informally, differential privacy is a property of an algorithm A, and states that the distribution of the output of A is insensitive to changes in one individual user's input. Formally, for every pair of inputs x, x' differing on at most one user's data, and every set of possible outputs $Z, \mathbb{P}[A(x) \in Z] \leq e^{\varepsilon} \cdot \mathbb{P}[A(x') \in Z]$. A locally differentially private algorithm is one in which each user *i* applies a private algorithm A_i only to their data.

Most local protocols are *non-interactive*: each user *i* sends a single message that is independent of all other messages. Non-interactive protocols can thus be written as $A(x_1, \ldots, x_n) = f(A_1(x_1), \ldots, A_n(x_n))$ for some function *f*, where each algorithm A_i satisfies ε -differential privacy. Our model requires an *interactive* protocol: each user *i* sends several messages over time, and these may depend on the messages sent by other users. This necessitates a slightly more complex formalism.

We consider interactive protocols among the *n* users and an additional center. Each user runs an algorithm A_i (possibly taking a private input x_i) and the central party runs an algorithm *C*. We let the random variable tr (A_1, \ldots, A_n, C) denote the *transcript* containing all the messages sent by all of the parties. For a given party *i* and a set of algorithms A'_{-i}, C' , we let tr_i $(x_i; A'_{-i}, C')$ denote the messages sent by user *i* in the transcript tr $(A_i(x_i), A'_{-i}, C')$. As a shorthand we will write tr_i (x_i) , since A'_{-i}, C' will be clear from context. We say that the protocol is locally differentially private if the function tr_i (x_i) is differentially private for every user *i* and every (possibly malicious) A'_{-i}, C' .

Definition 2.1. An interactive protocol (A_1, \ldots, A_n, C) satisfies ε -local differential privacy if for every user *i*, every pair of inputs x_i, x'_i for user *i*, and every set of algorithms A'_{-i}, C' , the resulting algorithm $\operatorname{tr}_i(x_i) = \operatorname{tr}_i(A_i(x_i), A'_{-i}, C')$ is ε -differentially private. That is, for every set of possible outputs Z, $\mathbb{P}[\operatorname{tr}_i(x_i) \in Z] \leq e^{\varepsilon} \cdot \mathbb{P}[\operatorname{tr}_i(x'_i) \in Z]$.

3 Overview: The THRESH Algorithm

Here we present our main algorithm, THRESH. The algorithmic framework is quite general, but for this high level overview we focus on the simplest setting where the data is Bernoulli. In Section 4 we formally present the algorithm for the Bernoulli case and analyze the algorithm to prove Theorem 1.1.

To explain the algorithm we first recall the distributional model. There are n users, each of whom belongs to a subgroup S_j for some $j \in [m]$; denote user *i*'s subgroup by g(i). There are $R = T\ell$ rounds divided into T epochs of length ℓ , denoted E^1, \ldots, E^T . In each round r, each user *i* receives a private bit $x_i^r \sim \text{Ber}(\mu_{g(i)}^r)$. We define the population-wide mean by $\mu^r = \frac{1}{n}(|S_1|\mu_1^r + \ldots + |S_m|\mu_m^r)$. For each epoch t, we use p^t to denote the average of the Bernoulli means during epoch t, $p^t = \frac{1}{\ell} \sum_{r \in E^t} \mu^r$. After every epoch t, our protocol outputs \tilde{p}^t such that $|p^t - \tilde{p}^t|$ is small.

The goal of THRESH is to maintain some public global estimate \tilde{p}^t of p^t . After any epoch t, we can *update* this global estimate \tilde{p}^t using *randomized response*: each user submits some differentially private estimate of the mean of their data, and the center aggregates these responses to obtain \tilde{p}^t . The main idea of THRESH is therefore to update the global estimate only when it might become sufficiently inaccurate, and thus take advantage of the possibly small number of changes in the underlying statistic p^t . The challenge is to privately identify when to update the global estimate.

The Voting Protocol. We identify these "update needed" epochs through a *voting protocol*. Users will examine their data and privately publish a vote for whether they believe the global estimate needs to be updated. If enough users vote to update the global estimate, we do so (using randomized

response). The challenge for the voting protocol is that users must use randomization in their voting process, to keep their data private, so we can only detect when a large number of users vote to update.

First, we describe a naïve voting protocol. In each epoch t, each user i computes a binary vote a_i^t . This vote is 1 if the user concludes from their own samples that the global estimate \tilde{p}^{t-1} is inaccurate, and 0 otherwise. Each user casts a noisy vote using randomized response accordingly, and if the sum of the noisy votes is large enough then a global update occurs.

The problem with this protocol is that small changes in the underlying mean p^t may cause some users to vote 1 and others to vote 0, and this might continue for an arbitrarily long time without inducing a global update. As a result, each voter "wastes" privacy in every epoch, which is what we wanted to avoid. We resolve this issue by having voters also estimate their *confidence* that a global update needs to occur, and vote proportionally. As a result, voters who have high confidence will lose more privacy per epoch (but the need for a global update will be detected quickly), while voters with low confidence will lose privacy more slowly (but may end up voting for many rounds).

In more detail, each user *i* decides their confidence level by comparing $|\hat{p}^t - \hat{p}_i^{f(t)}|$ —the difference between the local average of their data in the current epoch and their local average the last time a global update occurred—to a small set of discrete *thresholds*. Users with the highest confidence will vote in every epoch, whereas users with lower confidence will only vote in a small subset of the epochs. We construct these thresholds and subsets so that in expectation no user votes in more than a constant number of epochs before a global update occurs, and the amount of privacy each user loses from voting will not grow with the number of epochs required before an update occurs.

4 THRESH: The Bernoulli Case

4.1 The THRESH Algorithm (Bernoulli Case)

We now present pseudocode for the algorithm THRESH, including both the general framework as well as the specific voting and randomized response procedures. We emphasize that the algorithm only touches user data through the subroutines VOTE, and EST, each of which accesses data from a single user in at most two epochs. Thus, it is an online local protocol in which user *i*'s response in epoch *t* depends only on user *i*'s data from at most two epochs *t* and *t'* (and the global information that is viewable to all users). THRESH uses carefully chosen *thresholds* $\tau_b = 2(b+1)\sqrt{\ln(12nT/\delta)/2\ell}$ for $b = -1, 0, \ldots, \lfloor \log(T) \rfloor$ to discretize the confidence of each user; see Section 4.2 for details on this choice.

We begin with a privacy guarantee for THRESH. Our proof uses the standard analysis of the privacy properties of randomized response, combined with the fact that users have a cap on the number of updates that prevents the privacy loss from accumulating. We remark that our privacy proof *does not* depend on distributional assumptions, which are only used for the proof of accuracy. We sketch a proof here. A full proof appears in Section 6 of the Appendix.

Theorem 4.1. The protocol THRESH satisfies ε -local differential privacy (Definition 2.1)

Proof Sketch: Naïvely applying composition would yield a privacy parameter that scales with T. Instead, we will rely on our defined privacy "caps" c_i^V and c_i^E that limit the number of truthful votes and estimates each user sends. Intuitively, each user sends at most $O(\frac{\varepsilon}{a} + \frac{\varepsilon}{b})$ messages that depend on their private data, and the rest are sampled independently of their private data. Thus, we need only bound the privacy "cost" of each of these $O(\frac{\varepsilon}{a} + \frac{\varepsilon}{b})$ elements of a user's transcript coming from a different distribution and bound the sum of the costs by ε .

4.2 Accuracy Guarantee

Our accuracy theorem needs the following assumption on L, the size of the smallest subgroup, to guarantee that a global update occurs whenever any subgroup has all of its member users vote "yes".

Assumption 4.2. $L > \left(\frac{3}{\sqrt{2}} + \frac{\sqrt{32}}{\varepsilon}\right) \sqrt{n \ln(12mT/\delta)}.$

This brings us to our accuracy theorem, followed by a proof sketch (see Appendix 7 for full details). **Theorem 4.3.** Given number of users n, number of subgroups m, smallest subgroup size L, number of rounds R, privacy parameter ε , and chosen epoch length ℓ and number of epochs $T = R/\ell$, with

Algorithm 1 Global Algorithm: THRESH

Require: number of users n, number of epochs T, minimum subgroup size L, number of subgroups m, epoch length ℓ , privacy parameter ε , failure parameter δ 1: Initialize global estimate $\tilde{p}^0 \leftarrow -1$ 1: Initialize global estimate $p^{*} \leftarrow -1$ 2: Initialize vote privacy counters $c_{1}^{V}, \ldots, c_{n}^{V} \leftarrow 0, \ldots, 0$ 3: Initialize estimate privacy counters $c_{1}^{E}, \ldots, c_{n}^{E} \leftarrow 0, \ldots, 0$ 4: Initialize vote noise level $a \leftarrow \frac{4\sqrt{2n\ln(12mT/\delta)}}{L - \frac{3}{\sqrt{2}}\sqrt{n\ln(12mT/\delta)}}$ 5: Initialize estimate noise level $b \leftarrow \frac{\sqrt{2 \ln(12T/\delta)/2n}}{\log(T)\sqrt{\ln(12nT/\delta)/2\ell} - \sqrt{\ln(12T/\delta)/2n}}$ 6: for each epoch $t \in [T]$ do for each user $i \in [n]$ do 7: 8: User *i* publishes $a_i^t \leftarrow \text{VOTE}(i, t)$ 9: end for GlobalUpdate^t $\leftarrow \left(\frac{1}{n}\sum_{i=1}^{n}a_{i}^{t} > \frac{1}{e^{a}+1} + \sqrt{\frac{\ln(10T/\delta)}{2n}}\right)$ 10: if GlobalUpdate^t then 11: 12: $f(t) \leftarrow t$ 13: for each $i \in [n]$ do User *i* publishes $\tilde{p}_i^t \leftarrow \text{EST}(i, t)$ 14: 15: end for Aggregate user estimates into global estimate: $\tilde{p}^t \leftarrow \frac{1}{n} \sum_{i=1}^n \frac{\tilde{p}_i^t(e^{b}+1)-1}{e^{b}-1}$ 16: 17: else $f(t) \leftarrow f(t-1)$ 18: for each $i \in [n]$ do 19: 20: User *i* publishes $\tilde{p}_i^t \leftarrow \text{Ber}(\frac{1}{e^{b+1}})$ 21: end for $\tilde{p}^t \leftarrow \tilde{p}^{t-1}$ 22: end if 23: 24: Analyst publishes \tilde{p}^t 25: end for

Algorithm 2 Local Subroutine: VOTE

Require: user *i*, epoch *t* 1: Compute local estimate $\hat{p}_i^t \leftarrow \frac{1}{\ell} \sum_{r \in E^t} x_i^r$ 2: $b^* \leftarrow$ highest *b* such that $|\hat{p}_i^t - \hat{p}_i^{f(t)}| > \tau_b$ 3: VoteYes $_i^t \leftarrow (c_i^V < \varepsilon/4 \text{ and } 2^{\lfloor \log T \rfloor - b^*} \text{ divides } t)$ 4: **if** VoteYes $_i^t$ **then** 5: $c_i^V \leftarrow c_i^V + a$ 6: $a_i^t \leftarrow \text{Ber}(\frac{e^a}{e^{a+1}})$ 7: **else** 8: $a_i^t \leftarrow \text{Ber}(\frac{1}{e^{a+1}})$ 9: **end if** 10: Output a_i^t

Algorithm 3 Local Subroutine: EST

```
Require: user i, epoch t

1: SendEstimate<sup>t</sup><sub>i</sub> \leftarrow \{c_i^E < \varepsilon/4\}

2: if SendEstimate<sup>t</sup><sub>i</sub> then

3: c_i^E \leftarrow c_i^E + b

4: \tilde{p}_i^t \leftarrow \text{Ber}(\frac{1+\hat{p}_i^t(e^{b}-1)}{e^{b}+1})

5: else

6: \tilde{p}_i^t \leftarrow \text{Ber}(\frac{1}{e^{b}+1})

7: end if

8: Output \tilde{p}_i^t
```

probability at least $1 - \delta$, in every epoch $t \in [T]$ such that fewer than

$$\frac{\varepsilon}{4} \cdot \min\left(\frac{L}{8\sqrt{2n\ln(12mT/\delta)}} - 1, \frac{1}{\sqrt{2}}\left[\log(T)\sqrt{\frac{n}{\ell}} - 1\right]\right)$$

changes have occurred in epochs $1, 2, \ldots, t$, THRESH outputs \tilde{p}^t such that

$$|\tilde{p}^t - p^t| \le 4(\lfloor \log(T) \rfloor + 2)\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}.$$

Proof Sketch: We begin by proving correctness of the voting process. Lemma 7.1 guarantees that if every user decides that their subgroup distribution has not changed then a global update does not occur, while Lemma 7.2 guarantees that if every user in some subgroup decides that a change *has* occurred, then a global update occurs. By Lemma 7.3, for each user *i* the individual user estimates driving these voting decisions are themselves accurate to within $t_{\ell} = O(\sqrt{\ln(nT/\delta)/\ell})$ of the true $\mu_{g(i)}^t$. Finally, by Lemma 7.4 guarantees that if every user decides that a change has occurred, then a global update occurs a global estimate \tilde{p}^t that is within t_{ℓ} of the true p^t .

To reason about how distribution changes across multiple epochs affect THRESH, we use the preceding results to show that the number of global updates never exceeds the number of distribution changes (Lemma 7.6). A more granular guarantee then bounds the number of changes any user detects—and the number of times they vote accordingly—as a function of the number of distribution changes (Lemma 7.7). These results enable us, in Lemma 7.8, to show that each change increases a user's vote privacy cap c_i^V by at most 2 and estimate privacy cap c_i^E by at most 1.

Finally, recall that THRESH has each user *i* compare their current local estimate \hat{p}_i^t to their local estimate in the last global update, $\hat{p}_i^{f(t)}$, to decide how to vote, with higher *thresholds* for $|\hat{p}_i^t - \hat{p}_i^{f(t)}|$ increasing the likelihood of a "yes" vote. This implies that if every user in some subgroup computes a local estimate \hat{p}_i^t such that $|\hat{p}_i^t - \hat{p}_i^{f(t)}|$ exceeds the highest threshold, then every user sends a "yes" vote and a global update occurs, bringing with it the accuracy guarantee of Lemma 7.4. In turn, we conclude that $|\hat{p}_i^t - p_i^t|$ never exceeds the highest threshold, and our accuracy result follows.

We conclude this section with a few remarks about THRESH. First, while the provided guarantee depends on the number of changes of any size, one can easily modify THRESH to be robust to changes of size $\leq c$, paying and additive c term in the accuracy. Second, the accuracy's dependence on ℓ offers guidance for its selection: roughly, for desired accuracy α , one should set $\ell = 1/\alpha^2$. Finally, in practice one may want to periodically assess how many users have exhausted their privacy budgets, which we can achieve by extending the voting protocol to estimate the fraction of "live" users. We primarily view this as an implementation detail outside of the scope of the exact problem we study.

5 An Application to Heavy Hitters

We now use the methods developed above to obtain similar guarantees for a common problem in local differential privacy known as *heavy hitters*. In this problem each of n users has their own dictionary value $v \in D$ (e.g. their homepage), and an aggregator wants to learn the most frequently held dictionary values (e.g. the most common homepages), known as "heavy hitters", while satisfying local differential privacy for each user. The heavy hitters problem has attracted significant attention [20, 17, 5, 8]. Here, we show how our techniques combine with an approach of Bassily and Smith [3] to obtain the first guarantees for heavy hitters on evolving data. We note that our focus on this approach is primarily for expositional clarity; our techniques should apply just as well to other variants, which can lead to more efficient algorithms.

5.1 Setting Overview

As in the simpler Bernoulli case, we divide time into $\ell \cdot T$ rounds and T epochs. Here, in each round r each user i draws a sample v_i^r from a subgroup-specific distribution $\mathcal{P}_{g(i)}^r$ over the d values in dictionary \mathcal{D} and track $\mathcal{P}^1, \ldots, \mathcal{P}^T$, the weighted average dictionary distribution in each epoch. We will require the same Assumption 4.2 as in the Bernoulli case, and we also suppose that $d \gg n$, a common parameter regime for this problem.

In the Bernoulli case users could reason about the evolution of μ_j^t directly from their own ℓ samples in each epoch. Since it is reasonable to assume $d \gg \ell$, this is no longer possible in our new setting— \mathcal{P}_j^t is too large an object to estimate from ℓ samples. However, we can instead adopt a common approach in heavy hitters estimation and examine a "smaller" object using a *hash* on dictionary samples. We will therefore have users reason about the distribution p_j^t over hashes that \mathcal{P}_j^t induces, which is a much smaller joint distribution of m (transformed) Bernoulli distributions. Our hope is that users can reliably "detect changes" by analyzing p_j^t , and the feasibility of this method leans crucially on the properties of the hash in question.

5.2 Details and Privacy Guarantee

First we recall the details of the one-shot protocol from Bassily and Smith [3]. In their protocol, each user starts with a dictionary value $v \in [d]$ with an associated basis vector $e_v \in \mathbb{R}^d$. The user hashes this to a smaller vector $h \in \mathbb{R}^w$ using a (population-wide) Φ , a $w \times d$ Johnson-Lindenstrauss matrix where $w \ll d$. The user then passes this hash $\hat{z}_i^t = \Phi e_v$ to their own local randomizer \mathcal{R} , and the center aggregates these randomized values into a single \bar{z} which induces a frequency oracle.

We will modify this to produce a protocol HEAVYTHRESH in the vein of THRESH. In each epoch t each user i computes an estimated histogram \hat{p}_i^t and then hashes it into $\Phi \hat{p}_i^t \in \mathbb{R}^w$, where w = 20n (we assume the existence of a subroutine GenProj for generating Φ). Each user votes on whether or not a global update has occurred by comparing $\Phi \hat{p}_i^t$ to their estimate during the most recent update, $\Phi \hat{p}_i^{f(t)}$, in HEAVYVOTE. Next, HEAVYTHRESH aggregates these votes to determine whether or not a global update will occur. Depending on the result, each user then calls their own estimation subroutine HEAVYEST and outputs a randomized response using \mathcal{R} accordingly. If a global update occurs, HEAVYTHRESH aggregates these responses into a new published global hash \tilde{y}^t ; if not, HEAVYTHRESH publishes \tilde{y}^{t-1} . In either case, HEAVYTHRESH publishes (Φ, \tilde{y}^t) as well. This final output is a *frequency oracle*, which for any $v \in [d]$ offers an estimate $\langle \Phi e_v, \tilde{y}^t \rangle$ of $\mathcal{P}^t(v)$.

HEAVYTHRESH will use the following thresholds with $\tau_b = 2(b+1)\sqrt{2\ln(16wnT/\delta)/w\ell}$ for $b = -1, 0, \dots, \lfloor \log(T) \rfloor$. See Section 5.3 for details on this choice. Fortunately, the bulk of our analysis uses tools already developed either in Section 4 or Bassily and Smith [3]. Our privacy guarantee is almost immediate: since HEAVYTHRESH shares its voting protocols with THRESH, the only additional analysis needed is for the estimation randomizer \mathcal{R} (Lemma 8.1). Using the privacy of \mathcal{R} , privacy for HEAVYTHRESH follows by the same proof as for the Bernoulli case.

Theorem 5.1. HEAVYTHRESH is ϵ -local differentially private.

5.3 Accuracy Guarantee

As above, an accuracy guarantee for HEAVYTHRESH unfolds along similar lines as that for THRESH, with additional recourse to results from Bassily and Smith [3]. We again require Assumption 4.2 and also assume $d = 2^{o(n^2/\ell)}$ (a weak assumption made primarily for neatness in Theorem 1.2). Our result and its proof sketch follow, with details and full pseudocode in Appendix Section 9.

Theorem 5.2. With probability at least $1 - \delta$, in every epoch $t \in [T]$ such that fewer than

$$\frac{\varepsilon}{4} \cdot \min\left(\frac{L}{8\sqrt{2n\ln(12mT/\delta)}} - 1, \frac{\log(T)\sqrt{\frac{n\ln(320n^2T/\delta)}{10\ell}} - \sqrt{\frac{\ln(16dT/\delta)}{10}} - 2\ln(320nT/\delta)\sqrt{\frac{5}{n}}}{\sqrt{\ln(320nT/\delta)}}\right)$$

changes have occurred in epochs $1, 2, \ldots, t$,

$$|\hat{f}^{t}(v) - \mathcal{P}^{t}(v)| < 4(\log(T) + 2)\sqrt{\frac{2\ln(320n^{2}T/\delta)}{\ell}} + \sqrt{\frac{\ln(\frac{16ndT}{\delta})}{n}}$$

Proof Sketch: Our proof is similar to that of Theorem 4.3 and proceeds by proving analogous versions of the same lemmas, with users checking for changes in the subgroup distribution over observed *hashes* rather than observed *bits*. This leads to one new wrinkle in our argument: once we show that the globally estimated hash is close to the true hash, we must translate from closeness of hashes to closeness of the distributions they induce (Lemma 9.4). The rest of the proof, which uses guarantees of user estimate accuracy to 1. guarantee that sufficiently large changes cause global updates and 2. each change incurs a bounded privacy loss, largely follows that of Theorem 4.3.

References

- John M. Abowd. The challenge of scientific reproducibility and privacy protection for statistical agencies. Census Scientific Advisory Committee, 2016.
- [2] Differential Privacy Team Apple. Learning with privacy at scale. Technical report, Apple, 2017.
- [3] Raef Bassily and Adam Smith. Local, private, efficient protocols for succinct histograms. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 127–135. ACM, 2015.
- [4] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Differentially private empirical risk minimization: Efficient algorithms and tight error bounds. arXiv preprint arXiv:1405.7085, 2014.
- [5] Raef Bassily, Uri Stemmer, and Abhradeep Guha Thakurta. Practical locally private heavy hitters. In *Advances in Neural Information Processing Systems*, pages 2285–2293, 2017.
- [6] Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Usharsee Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. arXiv preprint arXiv:1710.00901, 2017.
- [7] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *Journal of the ACM (JACM)*, 60(2):12, 2013.
- [8] Mark Bun, Jelani Nelson, and Uri Stemmer. Heavy hitters and the structure of local privacy. *arXiv preprint arXiv:1711.04740*, 2017.
- [9] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109, 2011.
- [10] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In Proceedings of Advances in Neural Information Processing Systems 30 (NIPS 2017), 2017.
- [11] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science, 9(3–4):211–407, 2014.
- [12] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer, 2006.
- [13] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N Rothblum, and Salil Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 381–390. ACM, 2009.
- [14] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 715–724. ACM, 2010.
- [15] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 1054–1067. ACM, 2014.
- [16] Moritz Hardt and Guy N Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium* on, pages 61–70. IEEE, 2010.
- [17] Justin Hsu, Sanjeev Khanna, and Aaron Roth. Distributed private heavy hitters. In *International Colloquium on Automata, Languages, and Programming*, pages 461–472. Springer, 2012.
- [18] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? In *Proceedings of the 54th Annual Symposium on Foundations of Computer Science*, pages 531–540, 2008.

- [19] Katrina Ligett, Seth Neel, Aaron Roth, Bo Waggoner, and Steven Z Wu. Accuracy first: Selecting a differential privacy level for accuracy constrained erm. In Advances in Neural Information Processing Systems, pages 2563–2573, 2017.
- [20] Nina Mishra and Mark Sandler. Privacy via pseudorandom sketches. In Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems, pages 143–152. ACM, 2006.
- [21] Ryan M Rogers, Aaron Roth, Jonathan Ullman, and Salil Vadhan. Privacy odometers and filters: Pay-as-you-go composition. In Advances in Neural Information Processing Systems, pages 1921–1929, 2016.
- [22] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In Proceedings of the forty-second ACM symposium on Theory of computing, pages 765–774. ACM, 2010.
- [23] Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang. Privacy loss in apple's implementation of differential privacy on macos 10.12. *arXiv preprint arXiv:1709.02753*, 2017.
- [24] Jonathan Ullman. Tight lower bounds for locally differentially private selection. *Manuscript*, 2018.

6 Missing Proofs from Section 4

Theorem 6.1. The protocol THRESH satisfies ε -local differential privacy (Definition 2.1)

Proof. To begin, we fix an arbitrary private user i and arbitrary algorithms A'_{-i}, C' for the other users and for the center. Fix any pair of inputs x_i, x'_i for user i. To ease notation, let $\operatorname{tr} = \operatorname{tr}_i(A_i(x_i), A'_{-i}, C')$ and $\operatorname{tr}' = \operatorname{tr}_i(A_i(x'_i), A'_{-i}, C')$ be the random variables corresponding to the messages sent by user i in the protocol with inputs x_i, x'_i , respectively. Note that we drop the subscript i, since user i will be fixed throughout. To prove the theorem, it suffices to show

$$\frac{\mathbb{P}\left[\mathsf{tr}=z\right]}{\mathbb{P}\left[\mathsf{tr}'=z\right]} \le e^{\varepsilon}$$

for every possible set of messages z.

The structure of the transcripts is as follows: each epoch t contributes two elements, first the vote a^t (the output of VOTE(i,t)) and the estimate \tilde{p}^t (the output of EST(i,t)). So we can write $z = ((a^1, \tilde{p}^1), \dots, (a^T, \tilde{p}^T))$ and

$$\frac{\mathbb{P}\big[\mathsf{tr}=z\big]}{\mathbb{P}\big[\mathsf{tr}'=z\big]} = \prod_{t=1}^{T} \frac{\mathbb{P}\big[\mathsf{tr}'^{t}=(a^{t},\tilde{p}^{t}) \mid \mathsf{tr}^{< t}=z^{< t}\big]}{\mathbb{P}\big[\mathsf{tr}'^{t}=(a^{t},\tilde{p}^{t}) \mid \mathsf{tr}'^{< t}=z^{< t}\big]}.$$

For any execution of the protocol, we can partition the set of epochs into those epochs $S_V \subseteq [T]$ where in at least one of tr and tr' user *i* sets VoteYes_i^t to True, and those S_V^c where VoteYes_i^t is False in both tr and tr'; similarly, we can partition [T] into those epochs S_E where SendEstimate_i^t is True in at least one of tr and tr' and those S_E^c where SendEstimate_i^t is False in both tr and tr'.

Since every epoch in S_V causes the counter c_i^v to increase by a, S_V contains at most $\varepsilon/4a$ epochs from each of tr and tr', so $|S_V| \le \varepsilon/2a$.

For any $t \in S_V^c$, user *i* will sample a^t and \tilde{p}^t from $\text{Ber}(\frac{1}{e^{a+1}})$ in both tr and tr'. Thus

$$\prod_{t \in S_V^c} \frac{\mathbb{P}\left[\mathsf{tr}^t = (a^t, \tilde{p}^t) \mid \mathsf{tr}^{< t} = z^{< t}\right]}{\mathbb{P}\left[\mathsf{tr}'^t = (a^t, \tilde{p}^t) \mid \mathsf{tr}'^{< t} = z^{< t}\right]} = 1.$$

To complete the proof, we need to bound

$$\prod_{t \in S_V} \frac{\mathbb{P}\left[\mathsf{tr}^t = (a^t, \tilde{p}^t) \mid \mathsf{tr}^{$$

which will hold because every factor in the product is at most e^a and $|S_V| \le \varepsilon/2a$. To see why, consider some epoch $t \in S_V$. The first component of tr^t is the vote $a^t \in \{0, 1\}$. The only two possibilities for how a^t is chosen are $a^t \sim \operatorname{Ber}(\frac{1}{e^a+1})$ or $a^t \sim \operatorname{Ber}(\frac{e^a}{e^a+1})$. One can easily verify that for any $a^t \in \{0, 1\}$,

$$\frac{\mathbb{P}\left[a^t \mid \mathsf{tr}^{< t} = z^{< t}\right]}{\mathbb{P}\left[a^t \mid \mathsf{tr}^{\prime < t} = z^{< t}\right]} \le e^a.$$

We now consider the second component of tr^t , which is \tilde{p}^t . As in the S_V case, since every epoch in S_E causes the counter c_i^E to increase by b, S_E contains at most $\varepsilon/4b$ epochs from each of tr and tr', so $|S_E| \le \varepsilon/2b$.

When SendEstimate^t is False, then \tilde{p}^t is sampled from

$$\mathsf{Ber}\left(\frac{1}{e^b+1}\right)$$

and when SendEstimate^t is True, then \tilde{p}^t is sampled from

$$\mathsf{Ber}\left(\frac{1+\hat{p}^t(e^b-1)}{e^b+1}\right)$$

depending on the value of the private data \hat{p}^t , which lies in [0,1]. Thus, the parameter in the Bernoulli distribution lies in $\left[\frac{1}{e^{b+1}}, \frac{e^{b}}{e^{b+1}}\right]$. Again, one can easily verify that for any $\tilde{p}^t \in \{0, 1\}$,

$$\frac{\mathbb{P}\left[\tilde{p}^t \mid \mathsf{tr}^{< t} = z^{< t}, a^t\right]}{\mathbb{P}\left[\tilde{p}^t \mid \mathsf{tr}'^{< t} = z^{< t}, a^t\right]} \le e^b.$$

Putting it together, we have

$$\begin{split} \frac{\mathbb{P}\left[\mathsf{tr}=z\right]}{\mathbb{P}\left[\mathsf{tr}'=z\right]} &= \prod_{t=1}^{T} \frac{\mathbb{P}\left[\mathsf{tr}^{t}=\left(a^{t},\tilde{p}^{t}\right) \mid \mathsf{tr}^{$$

This completes the proof.

Missing Proofs From Section 4.2

We begin the proof of our accuracy guarantee with a series of lemmas. Recalling that we set

$$a = \frac{4\sqrt{2n\ln(12mT/\delta)}}{L - \frac{3}{\sqrt{2}}\sqrt{n\ln(12mT/\delta)}}$$

and

7

$$b = \frac{\sqrt{2\ln(12T/\delta)/2n}}{\log(T)\sqrt{\ln(12nT/\delta)/2\ell} - \sqrt{\ln(12T/\delta)/2n}}$$

we start by showing that if every user votes that a change has not occurred, then a global update will not occur.

Lemma 7.1. With probability at least $1 - \frac{\delta}{6}$, in every epoch $t \in [T]$, if every user *i* sets VoteYes_i^t \leftarrow False then GlobalUpdate^t \leftarrow False.

Proof. Since every user i sets VoteYes^t_i \leftarrow False, every a_i^t is an iid draw from a Bern $\left(\frac{1}{e^{a+1}}\right)$ distribution. Thus a Chernoff bound says

$$\mathbb{P}\left[\left|\frac{1}{n}\sum_{i=1}^{n}a_{i}^{t}-\frac{1}{e^{a}+1}\right| \geq \sqrt{\frac{\ln(12T/\delta)}{2n}}\right] \leq \frac{\delta}{6T}.$$

	٦	
	1	
	1	

Since GlobalUpdate^t $\leftarrow \left(\frac{1}{n}\sum_{i=1}^{n}a_{i}^{t} > \frac{1}{e^{a}+1} + \sqrt{\frac{\ln(12T/\delta)}{2n}}\right)$, GlobalUpdate^t \leftarrow False. Union-bounding across T epochs completes the proof

Similarly, we also want to ensure that if every user in some subgroup votes that an update has occurred then a global update will indeed occur.

Lemma 7.2. With probability at least $1 - \frac{\delta}{3}$, in every epoch $t \in [T]$, if there is a subgroup j where every user $i \in S_j$ sets VoteYes^t_i \leftarrow True then GlobalUpdate^t \leftarrow True.

Proof. Since $|S_j| \ge L$, Chernoff bounds imply that the aggregate vote satisfies

$$\frac{1}{n}\sum_{i=1}^{n}a_{i}^{t} > \frac{1}{n}\left[\frac{Le^{a}}{e^{a}+1} - \sqrt{\frac{L\ln(12mT/\delta)}{2}} + \frac{n-L}{e^{a}+1} - \sqrt{\frac{(n-L)\ln(12mT/\delta)}{2}}\right]$$

Recalling that GlobalUpdate^t \leftarrow True if and only if $\frac{1}{n}\sum_{i=1}^{n}a_{i}^{t} > \frac{1}{e^{a}+1} + \sqrt{\frac{\ln(12T/\delta)}{2n}}$, it suffices to show that

$$\frac{1}{n} \left[\frac{Le^a}{e^a + 1} - \sqrt{\frac{L\ln(12mT/\delta)}{2}} + \frac{n - L}{e^a + 1} - \sqrt{\frac{(n - L)\ln(12mT/\delta)}{2}} \right] > \frac{1}{e^a + 1} + \sqrt{\frac{\ln(12T/\delta)}{2n}}.$$

Rearranging, it is enough to show that

$$L > \frac{3}{\sqrt{2}} \cdot \frac{e^a + 1}{e^a - 1} \cdot \sqrt{n \ln(12mT/\delta)}$$

and using the fact that $\frac{e^a+1}{e^a-1} < \frac{a+2}{a}$ it is enough that

$$a > \frac{3\sqrt{2n\ln(12mT/\delta)}}{L - \frac{3}{\sqrt{2}}\sqrt{n\ln(12mT/\delta)}}$$

which follows from our setting of a. Union-bounding across m subgroups and T epochs completes the proof.

We now show that every user in every epoch obtains an estimate \hat{p}_i^t of $\mu_{g(i)}^t$ of bounded inaccuracy. This will enable us to-among other things-guarantee that users do not send "false positive" votes.

Lemma 7.3. With probability at least $1 - \frac{\delta}{6}$, in each epoch $t \in [T]$ each user *i* has

$$|\hat{p}_i^t - \mu_{g(i)}^t| < \sqrt{\frac{\ln(12nT/\delta)}{2\ell}}.$$

Proof. $\mathbb{E}\left[\hat{p}_{i}^{t}\right] = \mu_{q(i)}^{t}$, so by an additive Chernoff bound

$$\mathbb{P}\left[|\hat{p}_i^t - \mu_{g(i)}^t| \ge \sqrt{\frac{\ln(12nT/\delta)}{2\ell}}\right] \le 2\exp\left(-2\left[\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}\right]^2 \ell\right) = \delta/6nT.$$

A union bound across n users and T epochs then completes the proof.

Next, in those epochs in which a global update occurs and no user i has hit their estimation privacy cap c_i^E , in the interest of asymptotic optimality we want to obtain a similar error for the resulting collated estimate \tilde{p}^t .

Lemma 7.4. With probability at least $1 - \frac{\delta}{3}$, in every epoch $t \in [T]$ where every user i sets SendEstimate^t_i \leftarrow True,

$$\left|p^{t}-\tilde{p}^{t}\right| < 2(\log(T)+2)\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}.$$

Proof. Since every user *i* sets SendEstimate^{*t*}_{*i*} \leftarrow True we know that for all *i*

$$\tilde{p}_i^t \sim \mathsf{Ber}\left(\frac{1+\hat{p}_i^t(e^b-1)}{e^b+1}\right),$$

so

$$\mathbb{E}\left[\tilde{p}^{t}\right] = \mathbb{E}\left[\frac{1}{n}\sum_{i=1}^{n}\frac{\tilde{p}_{i}^{t}(e^{b}+1)-1}{e^{b}-1}\right] = \frac{1}{n}\sum_{i=1}^{n}\frac{\mathbb{E}\left[\tilde{p}_{i}^{t}\right](e^{b}+1)-1}{e^{b}-1} = \frac{1}{n}\sum_{i=1}^{n}\hat{p}_{i}^{t}$$

Since \tilde{p}^t is an average of $\{\frac{-1}{e^{b}-1}, \frac{e^b}{e^{b}-1}\}$ -valued random variables, we transform it into the $\{0, 1\}$ -valued random variable

$$Y = \left(\tilde{p}^{t} + \frac{1}{e^{b} - 1}\right) \cdot \frac{e^{b} - 1}{e^{b} + 1}.$$

Applying an additive Chernoff bound as above yields

$$\mathbb{P}\left[\left|Y - \mathbb{E}\left[Y\right]\right| \ge \sqrt{\frac{\ln(12T/\delta)}{2n}}\right] \le \frac{\delta}{6T}$$

which implies that

$$\mathbb{P}\left[\left|\tilde{p}^{t} - \frac{1}{n}\sum_{i=1}^{n} \hat{p}_{i}^{t}\right| \geq \left(\frac{e^{b}+1}{e^{b}-1}\right)\sqrt{\frac{\ln(12T/\delta)}{2n}}\right] \leq \frac{\delta}{6T}.$$

Similarly, as $\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^{n}\hat{p}_{i}^{t}\right] = p^{t}$,

$$\mathbb{P}\left[\left|\frac{1}{n}\sum_{i=1}^{n}\hat{p}_{i}^{t}-p^{t}\right|\geq\sqrt{\frac{\ln(12T/\delta)}{2n}}\right]\leq\frac{\delta}{6T}$$

Combining these results in the triangle inequality yields that with probability at least $1 - \frac{\delta}{6T}$

$$|\tilde{p}^t - p^t| < 2\left(\frac{e^b + 1}{e^b - 1}\right)\sqrt{\frac{\ln(12T/\delta)}{2n}}$$

Since $\frac{e^{b}+1}{e^{b}-1} < \frac{b+2}{b}$, this implies that

$$|\tilde{p}^t - p^t| < 2\left(\frac{b+2}{b}\right)\sqrt{\frac{\ln(12T/\delta)}{2n}}$$

so to get

$$|\tilde{p}^t - p^t| < 2(\log(T) + 2)\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}$$

it is enough for

$$b > \frac{\sqrt{2\ln(12T/\delta)/n}}{(\log(T) + 2)\sqrt{\ln(12nT/\delta)/2\ell} - \sqrt{\ln(12T/\delta)/2n}}$$

Substituting in our setting of

$$b = \frac{\sqrt{2\ln(12T/\delta)/2n}}{\log(T)\sqrt{\ln(12nT/\delta)/2\ell} - \sqrt{\ln(12T/\delta)/2n}}$$

and union-bounding over T epochs completes the proof.

Finally, we use the above lemmas to reason about how long users' privacy budgets last. We'll first define a useful term for this claim.

Definition 7.5. We say a change Δ_t occurs in epoch $t \in [T]$ if there exists subgroup j such that $\mu_j^t \neq \mu_j^{t-1}$. Given changes Δ_{t_1} and Δ_{t_2} where $t_1 < t_2$, we say that Δ_{t_1} and Δ_{t_2} are adjacent changes if there does not exist a change Δ_{t_3} such that $t_1 < t_3 < t_2$.

This lets us prove the following lemma bounding the frequency of global updates.

Lemma 7.6. With probability at least $1 - \delta$, given adjacent changes Δ_{t_1} and Δ_{t_2} , GlobalUpdate^t \leftarrow True in at most one epoch $t \in [t_1, t_2)$.

Proof. First, with probability at least $1 - \delta$ all of the preceding lemma in this section apply, and we condition on them for the remainder of this proof.

Assume instead that GlobalUpdate^t \leftarrow True and GlobalUpdate^{t'} \leftarrow True as well for $t_1 \le t < t' \le t_2 - 1$, and that GlobalUpdate^{t₃} \leftarrow False for all $t_3 \in (t, t')$. Recall that by Lemma 7.1, if in epoch t' every user *i* sets VoteYes^{t'}_i \leftarrow False then

$$\frac{1}{n}\sum_{i=1}^{n}a_{i}^{t'} \le \frac{1}{e^{a}+1} + \sqrt{\frac{\ln(12T/\delta)}{2n}}$$

which means GlobalUpdate^{t'} \leftarrow False. Therefore since we know GlobalUpdate^{t'} \leftarrow True, it follows that at least one user i sets VoteYes^{t'} \leftarrow True. By the thresholding structure of THRESH, this implies that

$$|\hat{p}_{i}^{t'} - \hat{p}_{i}^{t}| > 2\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}.$$

Since Lemma 7.3 guarantees that both $\hat{p}_i^{t'}$ and \hat{p}_i^t are within $\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}$ of $\mu_{g(i)}^{t'}$ and $\mu_{g(i)}^t$ respectively, it follows that $\mu_{g(i)}^{t'} \neq \mu_{g(i)}^t$. This contradicts the fact that Δ_{t_1} and Δ_{t_2} were adjacent changes.

We similarly bound the frequency with which users vote that a change has occurred.

Lemma 7.7. With probability at least $1 - \delta$, given adjacent changes Δ_{t_1} and Δ_{t_2} such that a global update occurs in $t_3 \in [t_1, t_2)$, for each user *i* there is at most one epoch $t \in (t_3, t_2)$ where VoteYes^t_i \leftarrow True.

Proof. First, with probability at least $1 - \delta$ all of the preceding lemmas in this section apply, and we condition on them for the remainder of this proof. In particular, Lemma 7.6 implies that t_3 is the only epoch in $[t_1, t_2)$ in which a global update occurs.

For contradiction, let $t_3 < t_4 < t_5 < t_2$ and assume that user i sets VoteYes $_i^{t_4} \leftarrow$ True and VoteYes $_i^{t_5} \leftarrow$ True. Since there is only one $t \in [T-1]$ such that $2^{\lfloor \log(T) \rfloor - 1}$ divides t, and the construction of the Vote subroutine requires this for a user m to set VoteYes $_m^t \leftarrow$ True, without loss of generality we may suppose that $|\hat{p}_i^{t_4} - \hat{p}_i^{f(t_4)}| > T_b$ and $|\hat{p}_i^{t_5} - \hat{p}_i^{f(t_5)}| > T_{b'}$ where $b > b' \ge 1$. However, Lemma 7.3 then implies that every user m in $S_{g(i)}$ has $|\hat{p}_m^{t_5} - \hat{p}_i^{f(t_5)}| > T_b'$, so by Lemma 7.6 GlobalUpdate $_{t_5}^{t_5} \leftarrow$ True, a contradiction of t_3 being the only epoch in $[t_1, t_2)$ in which a global update occurs.

Our last lemma before our main theorem ties the above results together to relate changes to increases in users' privacy caps c_i . This will eventually let us lower bound the time for which THRESH outputs accurate results.

Lemma 7.8. Denote by c_i^t the value of c_i in epoch t. Then with probability at least $1 - \delta$, across all epochs, given any two adjacent changes Δ_{t_1} and Δ_{t_2} , for every user i

$$c_{i,E}^{t_2-1} \le c_{i,E}^{t_1-1} + 1$$
$$c_{i,V}^{t_2-1} \le c_{i,V}^{t_1-1} + 2.$$

and

Proof. First, with probability at least $1 - \delta$ all of the preceding lemmas in this section apply, and we condition on them for the remainder of this proof.

Fix a user *i*. First, $c_{i,E}$ increases in any epoch *t* where SendEstimate^{*t*}_{*i*} \leftarrow True. This only happens in epoch where GlobalUpdate^{*t*} \leftarrow True, and by Lemma 7.6, at most one global update occurs in epochs in $[t_1, t_2)$, so $c_{i,E}^{t_2-1} \le c_{i,E}^{t_1-1} + 1$. We analyze $c_{i,V}$ in two cases.

<u>Case 1</u>: For all $t \in [t_1, t_2)$, GlobalUpdate^t \leftarrow False. Here, since no global update occurs, if $c_{i,V}^{t_2-1} > c_i^{t_1} + 2$ then there must exist 3 epochs $t \in [t_1, t_2)$ where user *i* sets VoteYes_i^t \leftarrow True, a contradiction of Lemma 7.7.

<u>Case 2</u>: For some epoch $t^* \in [t_1, t_2)$, GlobalUpdate^{t^*} \leftarrow True. It then suffices to show that user i sets VoteYes^t_{$i} <math>\leftarrow$ True in at most two epochs $t \in [t_1, t_2 - 1]$ (possibly including t^*).</sub>

Assume instead that VoteYes^{t3}_i, VoteYes^{t4}_i, and VoteYes^{t5}_i \leftarrow True for distinct $t_3, t_4, t_5 \in [t_1, t_2)$. By Lemmas 7.3 and 7.4, VoteYes^t_i \leftarrow False in any epoch $t \in [t^* + 1, t_2)$. Therefore $t_3, t_4, t_5 \in [t_1, t^*]$, and at least two are in $[t_1, t^* - 1]$ and do not trigger a global update. This again contradicts Lemma 7.7. \Box

Taken together, these preliminary results let us prove our main accuracy theorem.

Theorem 7.9. With probability at least $1 - \delta$, in every epoch $t \in [T]$ such that fewer than

$$\frac{\varepsilon}{4} \cdot \min\left(\frac{L}{8\sqrt{2n\ln(12mT/\delta)}} - 1, \frac{1}{\sqrt{2}}\left[\log(T)\sqrt{\frac{n}{\ell}} - 1\right]\right)$$

changes have occurred in epochs $1, 2, \ldots, t$,

$$|\tilde{p}^t - p^t| \le 4(\lfloor \log(T) \rfloor + 2)\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}$$

Proof. First, with probability at least $1 - \delta$ all of the preceding lemmas and corollaries in this section apply, and we condition on them for the remainder of this proof. In particular, since

$$\min\left(\frac{\varepsilon}{8a}, \frac{\varepsilon}{4b}\right) = \frac{\varepsilon}{4} \cdot \min\left(\frac{1}{2a}, \frac{1}{b}\right)$$
$$= \frac{\varepsilon}{4} \cdot \min\left(\frac{L - \frac{3}{\sqrt{2}}\sqrt{n\ln(12mT/\delta)}}{8\sqrt{2n\ln(12mT/\delta)}}, \frac{\log(T)\sqrt{\ln(12nT/\delta)/2\ell} - \sqrt{\ln(12T/\delta)/2n}}{\sqrt{2\ln(12T/\delta)/2n}}\right)$$
$$> \frac{\varepsilon}{4} \cdot \min\left(\frac{L}{8\sqrt{2n\ln(12mT/\delta)}} - 1, \frac{1}{\sqrt{2}}\left[\log(T)\sqrt{\frac{n}{\ell}} - 1\right]\right)$$

we know that the number of changes thus far is less than $\min\left(\frac{\varepsilon}{8a}, \frac{\varepsilon}{4b}\right)$, and by Lemma 7.8 for every user $i, c_i^V < \varepsilon/4$ and $c_i^E < \varepsilon/4$, i.e. no user has exceeded their voting or estimation privacy budget. Now suppose for contradiction that

$$|\tilde{p}^t - p^t| > 4(\lfloor \log(T) \rfloor + 2) \sqrt{\frac{\ln(12nT/\delta)}{2\ell}}$$

By Lemma 7.4 this means GlobalUpdate^t \leftarrow False, so by Lemma 7.2 for every subgroup $j \in [m]$ there exists user $i \in S_j$ such that

$$|\hat{p}_i^t - \hat{p}_i^{f(t)}| \le 2(\lfloor \log(T) \rfloor + 1)\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}$$

Lemma 7.3 then implies that

$$|\mu_j^t - \mu_j^{f(t)}| \le 2(\lfloor \log(T) \rfloor + 2)\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}.$$

Since this holds for every subgroup j, we get that

$$|p^t - p^{f(t)}| \le 2(\lfloor \log(T) \rfloor + 2)\sqrt{\frac{\ln(12nT/\delta)}{2\ell}}$$

and since GlobalUpdate^t \leftarrow False, by Lemma 7.4 this means that $\tilde{p}^t = \tilde{p}^{f(t)}$ and

$$|\tilde{p}^t - p^t| \le 4(\lfloor \log(T) \rfloor + 2) \sqrt{\frac{\ln(12nT/\delta)}{2\ell}}.$$

a contradiction.

8 Missing Proofs from Section 5.2

We start with full pseudocode for HEAVYTHRESH.

Next, we prove a lemma guaranteeing the privacy of the \mathcal{R} subroutine.

Lemma 8.1. \mathcal{R} is ε -locally DP.

Proof. Choose a possible output z of \mathcal{R} . Let h_1 and h_2 be two arbitrary input hashes. It suffices to show

$$\frac{\mathbb{P}\left[\mathcal{R}(h_1) = z\right]}{\mathbb{P}\left[\mathcal{R}(h_2) = z\right]} \le e^{\varepsilon}$$

<u>Case 1</u>: h_1 and h_2 are zero vectors. Then $\mathcal{R}(h_1)$ and $\mathcal{R}(h_2)$ have identical output distributions and the result is immediate.

<u>Case 2:</u> Exactly one (WLOG h_1) hash is a nonzero vector. Then

$$\mathbb{P}\left[\mathcal{R}(h_2)=z\right]=\frac{1}{2w}.$$

Similarly,

$$\mathbb{P}[\mathcal{R}(h_1) = z] \leq \frac{1}{w} \cdot \left(\frac{1}{2} + \frac{1}{2c_{\varepsilon}}\right).$$

Therefore

$$\frac{\mathbb{P}\left[\mathcal{R}(h_1)=z\right]}{\mathbb{P}\left[\mathcal{R}(h_2)=z\right]} \le 1 + \frac{1}{c_{\varepsilon}} = 1 + \frac{e^{\varepsilon}-1}{e^{\varepsilon}+1} \le e^{\varepsilon}.$$

<u>Case 3:</u> Neither h_1 nor h_2 is a zero vector. Then by the logic above

$$\frac{\mathbb{P}\left[\mathcal{R}(h_1)=z\right]}{\mathbb{P}\left[\mathcal{R}(h_2)=z\right]} \leq \frac{1+\frac{e^{\varepsilon}-1}{e^{\varepsilon}+1}}{1-\frac{e^{\varepsilon}-1}{e^{\varepsilon}+1}} = \frac{e^{\varepsilon}+1+e^{\varepsilon}-1}{e^{\varepsilon}+1-e^{\varepsilon}+1} = e^{\varepsilon}.$$

9	Missing	Proofs	From	Section	53
,	wiissing	1 1 0 0 1 5	L I OIII	Section	5.5

First, recall that we set

$$a = \frac{4\sqrt{2n\ln(12mT/\delta)}}{L - \frac{3}{\sqrt{2}}\sqrt{n\ln(12mT/\delta)}}$$

and

$$b = \frac{2\left(\sqrt{\frac{\ln(16wT/\delta)}{nw}} + \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}\right)}{2(\log(T) + 2)\sqrt{\frac{2\ln(16wT/\delta)}{w\ell}} - 2\sqrt{\frac{\ln(16dT/\delta)}{2wn}} - \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}}$$

We start with a result about \mathcal{R} .

Lemma 9.1. For all $\epsilon > 0$ and $x \in \left[-\frac{1}{\sqrt{w}}, \frac{1}{\sqrt{w}}\right]^w$, $\mathbb{E}\left[\mathcal{R}(x, \varepsilon)\right] = x$.

Proof. In the case where x = 0, we get

$$\mathbb{E}\left[\mathcal{R}(x)\right]_{j} = \frac{1}{w} \cdot \left(-\frac{c_{\varepsilon}\sqrt{w}}{2} + \frac{c_{\varepsilon}\sqrt{w}}{2}\right) = 0,$$

and for $x \neq 0$

$$\mathbb{E}\left[\mathcal{R}(x)\right]_{j} = \frac{1}{w}\left[\left(\frac{1}{2} + \frac{x_{j}\sqrt{w}}{2c_{\varepsilon}}\right)c_{\varepsilon}\sqrt{w} + \left(\frac{1}{2} - \frac{x_{j}\sqrt{w}}{2c_{\varepsilon}}\right)(-c_{\varepsilon}\sqrt{w})\right] = x_{j}.$$

Algorithm 4 Global Algorithm: HEAVYTHRESH

Require: number of users n, number of epochs T, minimum subgroup size L, number of subgroups *m*, epoch length ℓ , privacy parameter ε , failure parameter δ 1: Initialize global estimate $\tilde{y}^0 \leftarrow -1$ 2: Initialize update counters $c_1, \ldots, c_n \leftarrow 0, 0, \ldots, 0$ 3: Initialize vote noise level $a \leftarrow \frac{4\sqrt{2n\ln(12mT/\delta)}}{L - \frac{3}{\sqrt{2}}\sqrt{n\ln(12mT/\delta)}}$ 4: Initialize estimate noise level $b \leftarrow \frac{2\left(\sqrt{\frac{\ln(16wT/\delta)}{nw}} + \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}\right)}{2(\log(T)+2)\sqrt{\frac{2\ln(16wT/\delta)}{w\ell}} - 2\sqrt{\frac{\ln(16wT/\delta)\sqrt{w}}{2wn}} - \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}}$ 5: $w \leftarrow 20n$ 6: Initialize JL matrix $\Phi \leftarrow \text{GenProj}(w, d)$ 7: for each epoch $t \in [T]$ do 8: for each user $i \in [n]$ do 9: User *i* publishes $a_i^t \leftarrow \text{HEAVYVOTE}(i, t)$ 10: end for $\text{GlobalUpdate}^{t} \leftarrow \left(\frac{1}{n} \sum_{i=1}^{n} a_{i}^{t} > \frac{1}{e^{a} + 1} + \sqrt{\frac{\ln(16T/\delta)}{2n}}\right)$ 11: 12: if GlobalUpdate^t then 13: $f(t) \leftarrow t$ 14: for each $i \in [n]$ do User *i* publishes $\tilde{z}_i^t \leftarrow \text{HEAVYEST}(i, t)$ 15: end for 16: Aggregate user estimates into global estimate: $\tilde{y}^t \leftarrow \frac{1}{n} \sum_{i=1}^n \tilde{z}^t_i$ 17: 18: else $f(t) \leftarrow f(t-1)$ 19: for each $i \in [n]$ do 20: 21: User *i* publishes $\tilde{z}_i^t \leftarrow \mathcal{R}(0, b)$ 22: end for $\tilde{y}^t \leftarrow \tilde{y}^{t-1}$ 23: 24: end if 25: Analyst publishes \tilde{y}^t 26: Analyst publishes FO^t \leftarrow (Φ, \tilde{y}^t) 27: end for

Algorithm 5 Local Subroutine: HEAVYVOTE

Require: user *i*, epoch *t* 1: Compute local estimate $\hat{p}_i^t \leftarrow \frac{1}{\ell} \sum_{r=(t-1)\ell+1}^{t\ell} v_i^r$ 2: Compute local hash $\hat{y}_i^t \leftarrow \Phi \hat{p}_i^t$ 3: $b^* \leftarrow$ highest *b* such that $||\hat{y}_i^t - \hat{y}_i^{f(t)}||_{\infty} > \tau_b$ 4: VoteYes_{*i*} $\leftarrow (c_i^V < \varepsilon/4a \text{ and } 2^{\lfloor \log T \rfloor - b^*} \text{ divides } t)$ 5: **if** VoteYes_{*i*} **then** 6: $c_i^V \leftarrow c_i^V + a$ 7: $a_i^t \leftarrow \text{Ber}(\frac{e^a}{e^{a+1}})$ 8: **else** 9: $a_i^t \leftarrow \text{Ber}(\frac{1}{e^{a+1}})$ 10: **end if** 11: Output a_i^t

Algorithm 6 Frequency Oracle: A_{FO}

Require: Frequency oracle $(\Phi, \frac{1}{n} \sum_{i=1}^{n} z_i)$, dictionary value to be estimated $v \in [d]$ 1: Output $\hat{f}(v) = \langle \Phi e_v, \bar{z} \rangle$

Algorithm 7 Local Subroutine: HEAVYEST

Require: user *i*, epoch *t* 1: SendEstimate^{*t*}_{*i*} $\leftarrow \{c^E_i < \varepsilon/4b\}$ 2: **if** SendEstimate^{*t*}_{*i*} **then** 3: $c_i \leftarrow c^E_i + b$ 4: $\tilde{z}^t_i \leftarrow \mathcal{R}(\hat{y}^t_i, b)$ 5: **else** 6: $\tilde{z}^t_i \leftarrow \mathcal{R}(0, b)$ 7: **end if** 8: Output \tilde{z}^t_i

Algorithm 8 Client Randomizer: \mathcal{R}

Require: Hashed histogram $h = \Phi \hat{p}_i^t$, privacy parameter ε 1: Sample $j \in [w]$ uniformly at random 4: if $h \neq 0$ then Draw $x \sim \text{Uni}(0,1)$ 5: if $x < \frac{1}{2} + \frac{h_j \sqrt{w}}{2c_{\varepsilon}}$ then $z_j \leftarrow c_{\varepsilon} \sqrt{w}$ else 6: 7: 8: $z_j \leftarrow -c_{\varepsilon}\sqrt{w}$ 9: end if 10: 11: else 12: $z_j \leftarrow c_{\varepsilon} \sqrt{w}$ or $-c_{\varepsilon} \sqrt{w}$ u.a.r 13: end if 14: Output z

Lemmas 7.1 and 7.2, since they cover portions of the voting process shared between VOTE and HEAVYVOTE, apply here with only their failure probabilities changed to $\delta/8$ and $\delta/4$. We start with an analogue of Lemma 7.3.

Lemma 9.2. With probability at least $1 - \frac{\delta}{8}$, for every epoch t and user i, defining by $p_{g(i)}^t$ the *d*-dimensional vector where $p_{g(i)}^t(q) = \mathcal{P}_{g(i)}^t(q)$,

$$\left\|\Phi\hat{p}_{i}^{t}-\Phi p_{g(i)}^{t}\right\|_{\infty} < \sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}}.$$

Proof. $\Phi \hat{p}_i^t$ is a vector with entries in $\pm \frac{1}{\sqrt{w}}$, so setting $X = \frac{\sqrt{w} \left(\Phi \hat{p}_i^t + \frac{1}{\sqrt{w}}\right)}{2}$ we get $X \in [0,1]^m$ where each index X_j has $\mathbb{E}[X_j] = \frac{\sqrt{w} \left((\Phi p_{g(i)}^t)_j + \frac{1}{\sqrt{w}}\right)}{2}$. A Chernoff bound then says that, with probability at least $1 - \frac{\delta}{8}$, for every user i and every epoch t

$$\|X - \mathbb{E}[X_j]\|_{\infty} < \sqrt{\frac{\ln(16wnT/\delta)}{2\ell}}.$$

Scaling this result by $\frac{2}{\sqrt{w}}$ and transforming X back into $\Phi \hat{p}_i^t$ yields the claim.

This brings us to an analogue of Lemma 7.4.

Lemma 9.3. With probability at least $1 - \frac{3\delta}{8}$, for every epoch t where every user i sets SendEstimate^t_i \leftarrow True,

$$\left\|\tilde{y}^{t} - \Phi p^{t}\right\|_{\infty} < 2(\log(T) + 2)\sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}}$$

Proof. By Lemma 9.1, $\mathbb{E}\left[\tilde{y}^t\right] = \frac{1}{n}\sum_{i=1}^n \Phi \hat{p}_i^t$, and we want to begin by bounding $\|\tilde{y}^t - \frac{1}{n}\sum_{i=1}^n \Phi \hat{p}_i^t\|$. First, since each of the *n* random variables \hat{z}_i^t that make up $\tilde{y}^t = \frac{1}{n}\sum_{i=1}^n \hat{z}_i^t$ is a zero vector except for an independently randomly chosen index $s \in [w]$, for each $s \in [w]$ we can bound the number N_j^t of vectors \hat{z}_i^t that are nonzero in index s by an additive Chernoff bound:

$$\mathbb{P}\left[N_s^t > \frac{n}{w} + \sqrt{\frac{n\ln(8wT/\delta)}{2}}\right] \le \frac{\delta}{8wT}$$

Union bounding over w indices, since \tilde{y}^t is normalized by 1/n, we get that

$$\tilde{y}^t \in \left[-c_{\varepsilon} \sqrt{w} \left(\frac{1}{w} + \sqrt{\frac{\ln(8wT/\delta)}{2n}} \right), c_{\varepsilon} \sqrt{w} \left(\frac{1}{w} + \sqrt{\frac{\ln(8wT/\delta)}{2n}} \right) \right]^w$$

Thus scaling, applying a Chernoff bound to each index, then re-scaling and union bounding over all w indices and T epochs gives us that with probability at least $1 - \frac{\delta}{8}$ in every epoch t where GlobalUpdate^t \leftarrow True

$$\begin{split} \left\| \tilde{y}^t - \frac{1}{n} \sum_{i=1}^n \Phi \hat{p}_i^t \right\|_{\infty} &< 2c_{\varepsilon} \left(\frac{1}{w} + \sqrt{\frac{\ln(8wT/\delta)}{2n}} \right) \sqrt{\frac{w \ln(16wT/\delta)}{2n}} \\ &< c_{\varepsilon} \left(\sqrt{\frac{2\ln(16wT/\delta)}{wn}} + \frac{\ln(16T/\delta)\sqrt{w}}{n} \right). \end{split}$$

Similarly,

$$\begin{split} \left\| \Phi p^t - \frac{1}{n} \sum_{i=1}^n \Phi \hat{p}_i^t \right\|_{\infty} &= \| \Phi \|_{\infty} \cdot \left\| p^t - \frac{1}{n} \sum_{i=1}^n \hat{p}_i^t \right\|_{\infty} \\ &\leq \frac{1}{\sqrt{w}} \cdot \sqrt{\frac{\ln(16dT/\delta)}{2n}} \\ &= \sqrt{\frac{\ln(16dT/\delta)}{2wn}} \end{split}$$

where the inequality holds with probability at least $1 - \frac{\delta}{8}$ in every epoch t by the definition of Φ and a Chernoff bound on the sampling error of n samples from \mathcal{P}^t , union bounded over d dictionary elements and T epochs. By triangle inequality,

$$\left\|\tilde{y}^{t} - \Phi p^{t}\right\|_{\infty} < c_{\varepsilon} \left(\sqrt{\frac{\ln(16wT/\delta)}{nw}} + \frac{\ln(16wT/\delta)\sqrt{w}}{n}\right) + \sqrt{\frac{\ln(16dT/\delta)}{2wn}}$$

and since $c_{\varepsilon} = \frac{e^{b}+1}{e^{b}-1} < \frac{b+2}{b}$, it is enough to set b such that

$$\frac{b+2}{b} \left(\sqrt{\frac{\ln(16wT/\delta)}{nw}} + \frac{\ln(16wT/\delta)\sqrt{w}}{n} \right) + \sqrt{\frac{\ln(16dT/\delta)}{2wn}} \le 2(\log(T)+2)\sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}}$$

Substituting in our value

$$b = \frac{2\left(\sqrt{\frac{\ln(16wT/\delta)}{nw}} + \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}\right)}{2(\log(T) + 2)\sqrt{\frac{2\ln(16wT/\delta)}{w\ell}} - 2\sqrt{\frac{\ln(16dT/\delta)}{2wn}} - \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}}$$

n.

yields the claim.

We'll need the following result to translate bounds on $\|\tilde{y}^t - \Phi p^t\|$ into accuracy bounds relative to \mathcal{P}^t . Lemma 9.4. With probability at least $1 - \frac{\delta}{8}$, in every epoch t if

$$\left\|\tilde{y}^t - \Phi p^t\right\|_{\infty} < B$$

then, denoting by \hat{f}^t the frequency oracle induced by (Φ, \tilde{y}^t) ,

$$\max_{v \in [d]} \left| \hat{f}^t(v) - \mathcal{P}^t(v) \right| \le B\sqrt{w} + 2\sqrt{\frac{\ln(16ndT/\delta)}{n}}$$

Proof. The outline of our proof is similar (and in some steps identical) to that of Theorem 2.5 in Bassily and Smith [3], but we provide it here for completeness. First,

$$\begin{aligned} \max_{v \in [d]} |\hat{f}^{t}(v) - \mathcal{P}^{t}(v)| &= \max_{v \in [d]} \left| \langle \tilde{y}^{t}, \Phi e_{v} \rangle - \langle p^{t}, e_{v} \rangle \right| \\ &= \max_{v \in [d]} \left| \langle \tilde{y}^{t} - \Phi p^{t}, \Phi e_{v} \rangle + \langle \Phi p^{t}, \Phi e_{v} \rangle - \langle p^{t}, e_{v} \rangle \right| \\ &\leq \max_{v \in [d]} \left| \langle \tilde{y}^{t} - \Phi p^{t}, \Phi e_{v} \rangle \right| + \max_{v \in [d]} \left| \langle \Phi p^{t}, \Phi e_{v} \rangle - \langle p^{t}, e_{v} \rangle \right|. \end{aligned}$$

Then by Corollary 6.4 from Blum et al. [7], since $\gamma = \sqrt{\frac{\ln(16ndT/\delta)}{n}}$ and $w = 20n = \frac{20\ln(16ndT/\delta)}{\gamma^2}$, we get that with probability at least $1 - \frac{\delta}{8T}$ we can bound the second term by

$$\max_{v \in [d]} \left| \langle \Phi p^t, \Phi e_v \rangle - \langle p^t, e_v \rangle \right| \le \gamma \left(\left\| p^t \right\|_2^2 + \left\| e_v \right\|_2^2 \right) \le 2\gamma.$$

By our assumption we can bound the first term by

$$\left\|\tilde{y}^t - \Phi p^t\right\|_{\infty} \cdot \left\|\Phi e_v\right\|_1 \le B\sqrt{w}$$

Together with a union bound over the T epochs, these yield the claim.

We can use these lemmas to prove an analogue of Corollary 7.6. First, we specify our setting-specific redefinition of *change*.

Definition 9.5. We say a change Δ_t occurs in epoch $t \in [T]$ if there exists subgroup $j \in [m]$ such that $\mathcal{P}_j^t \neq \mathcal{P}_j^{t-1}$.

This lets us state the necessary result.

Lemma 9.6. With probability at least $1 - \delta$, given adjacent changes Δ_{t_1} and Δ_{t_2} , GlobalUpdate^t \leftarrow True in at most one epoch $t \in [t_1, t_2)$.

Proof. The proof is identical to that of Lemma 7.6, only replacing Lemma 7.3 with Lemma 9.2. \Box

Lemmas 7.7 and 7.8 apply in this setting unmodified, which finally lets us prove the following accuracy guarantee.

Theorem 9.7 (Accuracy Guarantee). With probability at least $1 - \delta$, in every epoch $t \in [T]$ such that fewer than

$$\frac{\varepsilon}{4} \cdot \min\left(\frac{L}{8\sqrt{2n\ln(12mT/\delta)}} - 1, \frac{\log(T)\sqrt{\frac{n\ln(320n^2T/\delta)}{10\ell}} - \sqrt{\frac{\ln(16dT/\delta)}{10}} - 2\ln(320nT/\delta)\sqrt{\frac{5}{n}}}{\sqrt{\ln(320nT/\delta)}\left(1 + \frac{20}{\sqrt{n}}\right)}\right)$$

changes have occurred in epochs $1, 2, \ldots, t$,

$$|\hat{f}^t(v) - \mathcal{P}^t(v)| < 4(\log(T) + 2)\sqrt{\frac{2\ln(320n^2T/\delta)}{\ell}} + \sqrt{\frac{\ln(\frac{16ndT}{\delta})}{n}}.$$

Proof. The proof is nearly identical to that of Theorem 4.3, replacing all lemmas with their heavyhitter analogues proven above. We provide it here for completeness. First, with probability at least $1 - \delta$ all of the preceding lemmas and corollaries in this section apply, and we condition on them for the remainder of this proof. In particular, since

$$\begin{split} \min\left(\frac{\varepsilon}{8a}, \frac{\varepsilon}{4b}\right) &= \frac{\varepsilon}{4} \cdot \min\left(\frac{1}{2a}, \frac{1}{b}\right) \\ &= \frac{\varepsilon}{4} \cdot \min\left(\frac{L - \frac{3}{\sqrt{2}}\sqrt{n\ln(12mT/\delta)}}{8\sqrt{2n\ln(12mT/\delta)}}, \frac{2(\log(T) + 2)\sqrt{\frac{2\ln(16wT/\delta)}{w\ell}} - 2\sqrt{\frac{\ln(16dT/\delta)}{2wn}} - \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}}{2\left(\sqrt{\frac{\ln(16wT/\delta)}{nw}} + \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}}\right) \right) \\ &> \frac{\varepsilon}{4} \cdot \min\left(\frac{L}{8\sqrt{2n\ln(12mT/\delta)}} - \frac{3}{16}, \frac{(\log(T) + 2)\sqrt{\frac{2\ln(320n^2T/\delta)}{20n\ell}} - \sqrt{\frac{\ln(16dT/\delta)}{40n^2}} - \frac{\ln(320nT/\delta)\sqrt{20n}}{n^2}}{\sqrt{\ln(320nT/\delta)}\left(\frac{1}{n} + \frac{20}{n^{3/2}}\right)}\right) \\ &= \frac{\varepsilon}{4} \cdot \min\left(\frac{L}{8\sqrt{2n\ln(12mT/\delta)}} - \frac{3}{16}, \frac{(\log(T) + 2)\sqrt{\frac{n\ln(320n^2T/\delta)}{10\ell}} - \frac{1}{2}\sqrt{\frac{\ln(16dT/\delta)}{10}} - 2\ln(320nT/\delta)\sqrt{\frac{5}{n}}}{\sqrt{\ln(320nT/\delta)}\left(1 + \frac{20}{\sqrt{n}}\right)}\right) \\ &> \frac{\varepsilon}{4} \cdot \min\left(\frac{L}{8\sqrt{2n\ln(12mT/\delta)}} - 1, \frac{\log(T)\sqrt{\frac{n\ln(320n^2T/\delta)}{10\ell}} - \sqrt{\frac{\ln(16dT/\delta)}{10}} - 2\ln(320nT/\delta)\sqrt{\frac{5}{n}}}{\sqrt{\ln(320nT/\delta)}\left(1 + \frac{20}{\sqrt{n}}\right)}\right) \end{split}$$

we know that the number of changes thus far is less than $\min\left(\frac{\varepsilon}{8a}, \frac{\varepsilon}{4b}\right)$, and by Lemma 7.8 for every user $i, c_i^V < \varepsilon/4$ and $c_i^E < \varepsilon/4$, i.e. no user has exceeded their voting or estimation privacy budget. Now suppose for contradiction that in epoch t

$$|\Phi p^t - \tilde{y}^{t-1}| > 4(\log(T) + 2)\sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}}.$$

By Lemma 9.3 this means GlobalUpdate^t \leftarrow False, so by Lemma 7.2 for every subgroup $j \in [m]$ there exists user $i \in S_j$ such that

$$\|\Phi \hat{p}_i^t - \Phi \hat{p}_i^{f(t)}\|_{\infty} \le 2(\log(T) + 1)\sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}}$$

Lemma 9.2 then implies that

$$\|\Phi p_{g(i)}^{t} - \Phi p_{g(i)}^{f(t)}\|_{\infty} \le 2(\log(T) + 2)\sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}}$$

Since this holds for every subgroup j, we get that

$$\|\Phi p^t - \Phi p^{f(t)}\|_{\infty} \le 2(\log(T) + 2)\sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}},$$

and since GlobalUpdate^t \leftarrow False, by Lemma 9.3 this means that $\tilde{y}^t = \tilde{y}^{f(t)}$, so

$$\|\tilde{y}^t - \Phi p^t\|\|_{\infty} \le 2(\log(T) + 2)\sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}}.$$

Plugging this quantity into Lemma 9.4 as B gives that for all $v \in [d]$

$$|\hat{f}^t(v) - \mathcal{P}^t(v)| < 2(\log(T) + 2)\sqrt{\frac{2\ln(16wnT/\delta)}{\ell}} + \sqrt{\frac{\ln(\frac{16ndT}{\delta})}{n}}.$$

Substituting w = 20n yields the claim.