

## LOCAL DIFFERENTIAL PRIVACY FOR EVOLVING DATA

MATTHEW JOSEPH, AARON ROTH, JONATHAN ULLMAN, AND BO WAGGONER

Department of Computer and Information Science, University of Pennsylvania  
*e-mail address:* majos@cis.upenn.edu

Department of Computer and Information Science, University of Pennsylvania  
*e-mail address:* aaroth@cis.upenn.edu

Khoury College of Computer Sciences, Northeastern University  
*e-mail address:* jullman@ccs.neu.edu

Microsoft Research New York  
*e-mail address:* bwag@colorado.edu

**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 (Dwork et al., 2006) is moving from theory to practice, with notable deployments by Google (Bittau et al., 2017; Erlingsson, Pihur and Korolova, 2014), Apple (Differential Privacy Team, Apple, 2017), Microsoft (Ding, Kulkarni and Yekhanin, 2017), and the U.S. Census Bureau (Abowd, 2016). 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 are collected by a trusted aggregator who performs and publishes the results of a differentially private computation (Dwork and Roth, 2014). However, Google, Apple, and Microsoft have instead chosen to operate in the *local model* (Bittau et al., 2017; Differential

*Key words and phrases:* local differential privacy.

Previous versions of this paper appeared at TPDP 2018 and NeurIPS 2018.

Privacy Team, Apple, 2017; Ding, Kulkarni and Yekhanin, 2017; Erlingsson, Pihur and Korolova, 2014), where users individually randomize their data on their own devices and send it to a potentially untrusted aggregator for analysis (Kasiviswanathan et al., 2008). In addition, the academic literature has largely focused on algorithms for performing one-time computations, like estimating many statistical quantities (Blum, Ligett and Roth, 2013; Hardt and Rothblum, 2010; Roth and Roughgarden, 2010) or training a classifier (Bassily, Smith and Thakurta, 2014; Chaudhuri, Monteleoni and Sarwate, 2011; Kasiviswanathan et al., 2008). Industrial applications, however have focused on tracking statistics about a user population, like the set of most frequently used emojis or words (Differential Privacy Team, Apple, 2017). 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. Naively 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  $\epsilon$ . This naïve strategy is what is used in practice (Bittau et al., 2017; Differential Privacy Team, Apple, 2017; Erlingsson, Pihur and Korolova, 2014). As a result, Tang et al. (2017) discovered that the privacy parameters guaranteed by Apple’s implementation of differentially private data collection can become unreasonably large even in relatively short time periods.<sup>1</sup> Published research on Google and Microsoft’s deployments suggests that they encounter similar issues (Bittau et al., 2017; Ding, Kulkarni and Yekhanin, 2017; Erlingsson, Pihur and Korolova, 2014).

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* (Dwork and Roth, 2014; Dwork et al., 2009; Hardt and Rothblum, 2010; Roth and Roughgarden, 2010) 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 (Kasiviswanathan et al., 2008; Ullman, 2018).

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

---

<sup>1</sup>Although the value of  $\epsilon$  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 (Bittau et al., 2017; Differential Privacy Team, Apple, 2017; Erlingsson, Pihur and Korolova, 2014). Thus, they may still provide a strong, if heuristic, privacy guarantee.

*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  $\ell$  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  $\varepsilon$  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 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 changes<sup>2</sup>.

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

---

<sup>2</sup>We can achieve a dynamic, data-dependent privacy guarantee using the notion of *ex-post* differential privacy (Ligett et al., 2017), for example by using a so-called privacy odometer (Rogers et al., 2016).

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

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

where  $L$  is the smallest subgroup size,  $n$  is the number of users,  $\ell$  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  $\alpha$  we need, ignoring log factors,  $\ell \geq 1/\alpha^2$  and that fewer than  $\varepsilon\alpha\sqrt{n}$  changes have occurred. We emphasize that our algorithm satisfies  $\varepsilon$ -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 Bassily and Smith (2015) 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  $\varepsilon$ -differentially private local protocol that achieves the following guarantee: with probability at least  $1 - \delta$ , while the total number of elapsed epochs  $t$  in which 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,  $\ell$  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, Pihur and Korolova (2014) on RAPPOR (the first large-scale deployment of differential privacy in the local model). Erlingsson, Pihur and Korolova (2014) 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, Kulkarni and Yekhanin (2017) 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., 2010). The canonical problem for differential privacy under 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.)

Subsequent work (Erlingsson et al., 2019) also studies locally private collection of data that changes over time. Unlike our work, their model does not assume that data is generated by an underlying evolving distribution. They instead suppose that data is deterministic and changes infrequently, while we suppose that data (which may change frequently) comes from a distribution that changes infrequently.

The local model was originally introduced by Kasiviswanathan et al. (2008), 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 (Bassily and Smith, 2015; Bassily, Stemmer and Thakurta, 2017; Bun, Nelson and Stemmer, 2017; Differential Privacy Team, Apple, 2017; Hsu, Khanna and Roth, 2012) and practical works (Bittau et al., 2017; Differential Privacy Team, Apple, 2017; Erlingsson, Pihur and Korolova, 2014).

## 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  $A(x_1, \dots, x_n) = f(A_1(x_1), \dots, A_n(x_n))$  for some function  $f$ , where each algorithm  $A_i$  satisfies  $\varepsilon$ -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  $\text{tr}(A_1, \dots, A_n, C)$  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  $\text{tr}_i(x_i; A'_{-i}, C')$  denote the messages sent by user  $i$  in the transcript  $\text{tr}(A_i(x_i), A'_{-i}, C')$ . As a shorthand we will write  $\text{tr}_i(x_i)$ , since  $A'_{-i}, C'$  will be clear from context. We say that the protocol is locally differentially private if the function  $\text{tr}_i(x_i)$  is differentially private for every user  $i$  and every (possibly malicious)  $A'_{-i}, C'$ .

**Definition 2.1 .** *An interactive protocol  $(A_1, \dots, A_n, C)$  satisfies  $\varepsilon$ -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  $\text{tr}_i(x_i) = \text{tr}_i(A_i(x_i), A'_{-i}, C')$  is  $\varepsilon$ -differentially private. That is, for every set of possible outputs  $Z$ ,  $\mathbb{P}[\text{tr}_i(x_i) \in Z] \leq e^\varepsilon \cdot \mathbb{P}[\text{tr}_i(x'_i) \in Z]$ .

### 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  $\ell$ , denoted  $E^1, \dots, E^T$ . 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  $\mu^r = \frac{1}{n}(|S_1|\mu_1^r + \dots + |S_m|\mu_m^r)$ . 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 (Figure 1), 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 (Figure 2), 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  $b = -1, 0, \dots, \lfloor \log(T) \rfloor$  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 A of the Appendix.

**Theorem 4.1 .** *The protocol THRESH satisfies  $\varepsilon$ -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  $\varepsilon$ .  $\square$

**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 B 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  $\varepsilon$ , and chosen epoch length  $\ell$  and number of epochs*

---

**Algorithm 1** Global Algorithm: THRESH

---

**Require:** number of users  $n$ , number of epochs  $T$ , minimum subgroup size  $L$ , number of subgroups  $m$ , epoch length  $\ell$ , privacy parameter  $\varepsilon$ , failure parameter  $\delta$

- 1: Initialize global estimate  $\tilde{p}^0 \leftarrow -1$
- 2: Initialize vote privacy counters  $c_1^V, \dots, c_n^V \leftarrow 0, \dots, 0$
- 3: Initialize estimate privacy counters  $c_1^E, \dots, c_n^E \leftarrow 0, \dots, 0$
- 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**
- 7:   **for** each user  $i \in [n]$  **do**
- 8:     User  $i$  publishes  $a_i^t \leftarrow \text{VOTE}(i, t)$
- 9:   **end for**
- 10:   GlobalUpdate $^t \leftarrow \left( \frac{1}{n} \sum_{i=1}^n a_i^t > \frac{1}{e^a+1} + \sqrt{\frac{\ln(10T/\delta)}{2n}} \right)$
- 11:   **if** GlobalUpdate $^t$  **then**
- 12:      $f(t) \leftarrow t$
- 13:     **for** each  $i \in [n]$  **do**
- 14:       User  $i$  publishes  $\tilde{p}_i^t \leftarrow \text{EST}(i, t)$
- 15:     **end for**
- 16:     Aggregate user estimates into global estimate:  $\tilde{p}^t \leftarrow \frac{1}{n} \sum_{i=1}^n \frac{\tilde{p}_i^t(e^b+1)-1}{e^b-1}$
- 17:   **else**
- 18:      $f(t) \leftarrow f(t-1)$
- 19:     **for** each  $i \in [n]$  **do**
- 20:       User  $i$  publishes  $\tilde{p}_i^t \leftarrow \text{Ber}(\frac{1}{e^b+1})$
- 21:     **end for**
- 22:      $\tilde{p}^t \leftarrow \tilde{p}^{t-1}$
- 23:   **end if**
- 24:   Analyst publishes  $\tilde{p}^t$
- 25: **end for**

---

FIGURE 1. Pseudocode for THRESH.

$T = R/\ell$ , 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  $1, 2, \dots, t$ , THRESH outputs  $\tilde{p}^t$  such that

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

*Proof Sketch:* We begin by proving correctness of the voting process. Lemma B.1 guarantees that if every user decides that their subgroup distribution has not changed then

---

**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:  $\text{VoteYes}_i^t \leftarrow (c_i^V < \varepsilon/4 \text{ and } 2^{\lfloor \log T \rfloor - b^*} \text{ divides } t)$
  - 4: **if**  $\text{VoteYes}_i^t$  **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:  $\text{SendEstimate}_i^t \leftarrow \{c_i^E < \varepsilon/4\}$
  - 2: **if**  $\text{SendEstimate}_i^t$  **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  $\tilde{p}_i^t$
- 

FIGURE 2. Pseudocode for VOTE and EST.

a global update does not occur, while Lemma B.2 guarantees that if every user in some subgroup decides that a change *has* occurred, then a global update occurs. By Lemma B.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 B.4 guarantees that if every user decides that a change has occurred, then a global update occurs that produces a global estimate  $\tilde{p}^t$  that is within  $t_\ell$  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 B.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 B.7). These results enable us, in Lemma B.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 B.4. In turn, we conclude that  $|\tilde{p}^t - p^t|$  never exceeds the highest threshold, and our accuracy result follows.  $\square$

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 an additive  $c$  term in the accuracy. Second, the accuracy’s dependence on  $\ell$  offers guidance for its selection: roughly, for desired accuracy  $\alpha$ , 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  $v \in \mathcal{D}$  (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 (Bassily, Stemmer and Thakurta, 2017; Bun, Nelson and Stemmer, 2017; Hsu, Khanna and Roth, 2012; Mishra and Sandler, 2006). Here, we show how our techniques combine with an approach of Bassily and Smith (2015) 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  $\mathcal{P}^1, \dots, \mathcal{P}^T$ , 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  $\mu_j^t$  directly from their own  $\ell$  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  $\ell$  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 (2015). 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)  $\Phi$ , 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  $\Phi$ ). 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  $(\Phi, \tilde{y}^t)$  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 (2015). 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 C.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  $\epsilon$ -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 (2015). 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 D.

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

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

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

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

*Proof 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 D.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.  $\square$

## REFERENCES

- Abowd, John M.** 2016. “The Challenge of Scientific Reproducibility and Privacy Protection for Statistical Agencies.” U.S. Census Bureau. <https://www2.census.gov/cac/sac/meetings/2016-09/2016-abowd.pdf>.
- Bassily, Raef, Adam Smith, and Abhradeep Thakurta.** 2014. “Differentially private empirical risk minimization: Efficient algorithms and tight error bounds.” *Foundations of Computer Science (FOCS)*. <https://doi.org/10.1109/FOCS.2014.56>.
- Bassily, Raef, and Adam Smith.** 2015. “Local, private, efficient protocols for succinct histograms.” *Symposium on the Theory of Computing (STOC)*. <https://doi.org/10.1145/2746539.2746632>.
- Bassily, Raef, Uri Stemmer, and Abhradeep Guha Thakurta.** 2017. “Practical locally private heavy hitters.” *Advances in Neural Information Processing Systems*. <https://papers.nips.cc/paper/6823-practical-locally-private-heavy-hitters>.
- Bittau, Andrea, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld.** 2017. “Prochlo: Strong Privacy for Analytics in the Crowd.” *SOSP ’17*, 441–459. <https://doi.org/10.1145/3132747.3132769>.
- Blum, Avrim, Katrina Ligett, and Aaron Roth.** 2013. “A learning theory approach to noninteractive database privacy.” *Symposium on the Theory of Computing (STOC)*. <https://doi.org/10.1145/1374376.1374464>.
- Bun, Mark, Jelani Nelson, and Uri Stemmer.** 2017. “Heavy Hitters and the Structure of Local Privacy.” *Principles of Database Systems (PODS)*. <https://doi.org/10.1145/3196959.3196981>.
- Chaudhuri, Kamalika, Claire Monteleoni, and Anand D Sarwate.** 2011. “Differentially private empirical risk minimization.” *Journal of Machine Learning Research*. <https://doi.org/10.5555/1953048.2021036>.
- Differential Privacy Team, Apple.** 2017. “Learning With Privacy at Scale.” Apple. <https://machinelearning.apple.com/docs/learning-with-privacy-at-scale/appledifferentialprivacysystem.pdf> (accessed 2020-01-11).
- Ding, Bolin, Janardhan Kulkarni, and Sergey Yekhanin.** 2017. “Collecting Telemetry Data Privately.” *Neural Information Processing Systems (NIPS)*. <https://papers.nips.cc/paper/6948-collecting-telemetry-data-privately.pdf>.
- Dwork, Cynthia, and Aaron Roth.** 2014. “The algorithmic foundations of differential privacy.” *Foundations and Trends® in Theoretical Computer Science*, 9(3–4): 211–407. <https://www.cis.upenn.edu/~aaroth/privacybook.html>.
- Dwork, Cynthia, Frank McSherry, Kobbi Nissim, and Adam Smith.** 2006. “Calibrating noise to sensitivity in private data analysis.” *Theory of Cryptography Conference (TCC)*. [https://doi.org/10.1007/11681878\\_14](https://doi.org/10.1007/11681878_14).
- Dwork, Cynthia, Moni Naor, Omer Reingold, Guy N Rothblum, and Salil Vadhan.** 2009. “On the complexity of differentially private data release: efficient algorithms and hardness results.” *Symposium on the Theory of Computing (STOC)*. <https://doi.org/10.1145/1536414.1536467>.
- Dwork, Cynthia, Moni Naor, Toniann Pitassi, and Guy N Rothblum.** 2010. “Differential privacy under continual observation.” *Symposium on the Theory of Computing (STOC)*. <https://doi.org/10.1145/1806689.1806787>.

- Erlingsson, Úlfar, Vasyl Pihur, and Aleksandra Korolova.** 2014. “Rappor: Randomized aggregatable privacy-preserving ordinal response.” *Conference on Computer and Communications Security (CCS)*. <https://doi.org/10.1145/2660267.2660348>.
- Erlingsson, Úlfar, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta.** 2019. “Amplification by Shuffling: From Local to Central Differential Privacy via Anonymity.” *Symposium on Discrete Algorithms (SODA)*. <https://doi.org/10.5555/3310435.3310586>.
- Hardt, Moritz, and Guy N Rothblum.** 2010. “A multiplicative weights mechanism for privacy-preserving data analysis.” *Foundations of Computer Science (FOCS)*. <https://ieeexplore.ieee.org/document/5670948>.
- Hsu, Justin, Sanjeev Khanna, and Aaron Roth.** 2012. “Distributed private heavy hitters.” *International Colloquium on Automata, Languages, and Programming (ICALP)*. [https://doi.org/10.1007/978-3-642-31594-7\\_39](https://doi.org/10.1007/978-3-642-31594-7_39).
- Kasiviswanathan, Shiva Prasad, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith.** 2008. “What can we learn privately?” *Foundations of Computer Science (FOCS)*. <https://ieeexplore.ieee.org/document/4690986>.
- Ligett, Katrina, Seth Neel, Aaron Roth, Bo Waggoner, and Steven Z Wu.** 2017. “Accuracy First: Selecting a Differential Privacy Level for Accuracy Constrained ERM.” *Neural Information Processing Systems (NIPS)*. <https://papers.nips.cc/paper/6850-accuracy-first-selecting-a-differential-privacy-level-for-accuracy-constrained-erm>.
- Mishra, Nina, and Mark Sandler.** 2006. “Privacy via pseudorandom sketches.” *Principles of Database Systems (PODS)*. <https://doi.org/10.1145/1142351.1142373>.
- Rogers, Ryan M, Aaron Roth, Jonathan Ullman, and Salil Vadhan.** 2016. “Privacy odometers and filters: Pay-as-you-go composition.” *Neural Information Processing Systems (NIPS)*. <https://papers.nips.cc/paper/6170-privacy-odometers-and-filters-pay-as-you-go-composition>.
- Roth, Aaron, and Tim Roughgarden.** 2010. “Interactive privacy via the median mechanism.” *Symposium on the Theory of Computing (STOC)*. <https://doi.org/10.1145/1806689.1806794>.
- Tang, Jun, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang.** 2017. “Privacy Loss in Apple’s Implementation of Differential Privacy on Mac OS 10.12.” *arXiv preprint arXiv:1709.02753*. <https://arxiv.org/abs/1709.02753>.
- Ullman, Jonathan.** 2018. “Tight Lower Bounds for Locally Differentially Private Selection.” *arXiv preprint arXiv:1802.02638*. <https://arxiv.org/abs/1802.02638>.

## APPENDIX A. MISSING PROOFS FROM SECTION 4

**Theorem A.1 .** *The protocol THRESH satisfies  $\varepsilon$ -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  $\text{tr} = \text{tr}_i(A_i(x_i), A'_{-i}, C')$  and  $\text{tr}' = \text{tr}_i(A_i(x'_i), A'_{-i}, C')$  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

$$\frac{\mathbb{P}[\text{tr} = z]}{\mathbb{P}[\text{tr}' = z]} \leq e^\varepsilon$$

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  $\text{VOTE}(i, t)$ ) and the estimate  $\tilde{p}^t$  (the output of  $\text{EST}(i, t)$ ). So we can write  $z = ((a^1, \tilde{p}^1), \dots, (a^T, \tilde{p}^T))$  and

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

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  $\text{tr}$  and  $\text{tr}'$  user  $i$  sets  $\text{VoteYes}_i^t$  to True, and those  $S_V^c$  where  $\text{VoteYes}_i^t$  is False in both  $\text{tr}$  and  $\text{tr}'$ ; similarly, we can partition  $[T]$  into those epochs  $S_E$  where  $\text{SendEstimate}_i^t$  is True in at least one of  $\text{tr}$  and  $\text{tr}'$  and those  $S_E^c$  where  $\text{SendEstimate}_i^t$  is False in both  $\text{tr}$  and  $\text{tr}'$ .

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  $\text{tr}$  and  $\text{tr}'$ , so  $|S_V| \leq \varepsilon/2a$ .

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  $\text{tr}$  and  $\text{tr}'$ . Thus

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

To complete the proof, we need to bound

$$\prod_{t \in S_V} \frac{\mathbb{P}[\text{tr}^t = (a^t, \tilde{p}^t) \mid \text{tr}^{<t} = z^{<t}]}{\mathbb{P}[\text{tr}'^t = (a^t, \tilde{p}^t) \mid \text{tr}'^{<t} = z^{<t}]} \leq e^{\varepsilon/2},$$

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

$$\frac{\mathbb{P}[a^t \mid \text{tr}^{<t} = z^{<t}]}{\mathbb{P}[a^t \mid \text{tr}'^{<t} = z^{<t}]} \leq e^a.$$

We now consider the second component of  $\text{tr}^t$ , 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  $\text{tr}$  and  $\text{tr}'$ , so  $|S_E| \leq \varepsilon/2b$ .

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

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

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

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

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  $[\frac{1}{e^b+1}, \frac{e^b}{e^b+1}]$ . Again, one can easily verify that for any  $\tilde{p}^t \in \{0, 1\}$ ,

$$\frac{\mathbb{P}[\tilde{p}^t | \text{tr}^{<t} = z^{<t}, a^t]}{\mathbb{P}[\tilde{p}^t | \text{tr}'^{<t} = z^{<t}, a^t]} \leq e^b.$$

Putting it together, we have

$$\begin{aligned} \frac{\mathbb{P}[\text{tr} = z]}{\mathbb{P}[\text{tr}' = z]} &= \prod_{t=1}^T \frac{\mathbb{P}[\text{tr}^t = (a^t, \tilde{p}^t) | \text{tr}^{<t} = z^{<t}]}{\mathbb{P}[\text{tr}'^t = (a^t, \tilde{p}^t) | \text{tr}'^{<t} = z^{<t}]} \\ &= \prod_{t \in S_V} \frac{\mathbb{P}[\text{tr}^t = a^t | \text{tr}^{<t} = z^{<t}]}{\mathbb{P}[\text{tr}'^t = a^t | \text{tr}'^{<t} = z^{<t}]} \cdot \prod_{t \in S_E} \frac{\mathbb{P}[\text{tr}^t = \tilde{p}^t | \text{tr}^{<t} = z^{<t}, a^t]}{\mathbb{P}[\text{tr}'^t = \tilde{p}^t | \text{tr}'^{<t} = z^{<t}, a^t]} \\ &\leq e^{a \cdot |S_V|} \cdot e^{b \cdot |S_E|} \leq e^{\varepsilon/2} \cdot e^{\varepsilon/2} \leqq e^\varepsilon. \end{aligned}$$

This completes the proof.  $\square$

## APPENDIX B. 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

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

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

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

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

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

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

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 B.2 .** *With probability at least  $1 - \delta/3$ , in every epoch  $t \in [T]$ , if there is a subgroup  $j$  where every user  $i \in S_j$  sets  $\text{VoteYes}_i^t \leftarrow \text{True}$  then  $\text{GlobalUpdate}^t \leftarrow \text{True}$ .*

*Proof.* Since  $|S_j| \geq L$ , 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  $\text{GlobalUpdate}^t \leftarrow \text{True}$  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.  $\square$

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 B.3 .** *With probability at least  $1 - \delta/6$ , 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.*  $\mathbb{E}[\hat{p}_i^t] = \mu_{g(i)}^t$ , so by an additive Chernoff bound

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

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

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 B.4 .** *With probability at least  $1 - \frac{\delta}{3}$ , in every epoch  $t \in [T]$  where every user  $i$  sets  $\text{SendEstimate}_i^t \leftarrow \text{True}$ ,*

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

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

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

so

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

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

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

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  $\mathbb{E} [\frac{1}{n} \sum_{i=1}^n \hat{p}_i^t] = p^t$ ,

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

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 suffices that

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

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.  $\square$

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 B.5 .** *We say a change  $\Delta_t$  occurs in epoch  $t \in [T]$  if there exists subgroup  $j$  such that  $\mu_j^t \neq \mu_j^{t-1}$ . Given changes  $\Delta_{t_1}$  and  $\Delta_{t_2}$  where  $t_1 < t_2$ , we say that  $\Delta_{t_1}$  and  $\Delta_{t_2}$  are adjacent changes if there does not exist a change  $\Delta_{t_3}$  such that  $t_1 < t_3 < t_2$ .*

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

**Lemma B.6 .** *With probability at least  $1 - \delta$ , given adjacent changes  $\Delta_{t_1}$  and  $\Delta_{t_2}$ ,  $\text{GlobalUpdate}^t \leftarrow \text{True}$  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  $\text{GlobalUpdate}^t \leftarrow \text{True}$  and  $\text{GlobalUpdate}^{t'} \leftarrow \text{True}$  as well for  $t_1 \leq t < t' \leq t_2 - 1$ , and that  $\text{GlobalUpdate}^{t_3} \leftarrow \text{False}$  for all  $t_3 \in (t, t')$ . Recall that by Lemma B.1, if in epoch  $t'$  every user  $i$  sets  $\text{VoteYes}_i^{t'} \leftarrow \text{False}$  then

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

which means  $\text{GlobalUpdate}^{t'} \leftarrow \text{False}$ . Therefore since we know  $\text{GlobalUpdate}^{t'} \leftarrow \text{True}$ , it follows that at least one user  $i$  sets  $\text{VoteYes}_i^{t'} \leftarrow \text{True}$ . 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 B.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  $\Delta_{t_1}$  and  $\Delta_{t_2}$  were adjacent changes.  $\square$

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

**Lemma B.7 .** *With probability at least  $1 - \delta$ , given adjacent changes  $\Delta_{t_1}$  and  $\Delta_{t_2}$  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  $\text{VoteYes}_i^t \leftarrow \text{True}$ .*

*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 B.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  $\text{VoteYes}_i^{t_4} \leftarrow \text{True}$  and  $\text{VoteYes}_i^{t_5} \leftarrow \text{True}$ . 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  $\text{VoteYes}_m^t \leftarrow \text{True}$ , 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  $b > b' \geq 1$ . However, Lemma B.3 then implies that every user  $m$  in  $S_{g(i)}$  has  $|\hat{p}_m^{t_5} - \hat{p}_m^{f(t_5)}| > T'_b$ , so by Lemma B.6  $\text{GlobalUpdate}^{t_5} \leftarrow \text{True}$ , a contradiction of  $t_3$  being the only epoch in  $[t_1, t_2]$  in which a global update occurs.  $\square$

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 B.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  $\Delta_{t_1}$  and  $\Delta_{t_2}$ , for every user  $i$

$$c_{i,E}^{t_2-1} \leq c_{i,E}^{t_1-1} + 1$$

and

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

*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  $\text{SendEstimate}_i^t \leftarrow \text{True}$ . This only happens in epoch where  $\text{GlobalUpdate}^t \leftarrow \text{True}$ , and by Lemma B.6, at most one global update occurs in epochs in  $[t_1, t_2)$ , so  $c_{i,E}^{t_2-1} \leq c_{i,E}^{t_1-1} + 1$ . We analyze  $c_{i,V}$  in two cases.

Case 1: For all  $t \in [t_1, t_2)$ ,  $\text{GlobalUpdate}^t \leftarrow \text{False}$ . 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  $\text{VoteYes}_i^t \leftarrow \text{True}$ , a contradiction of Lemma B.7.

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

Assume instead that  $\text{VoteYes}_i^{t_3}$ ,  $\text{VoteYes}_i^{t_4}$ , and  $\text{VoteYes}_i^{t_5} \leftarrow \text{True}$  for distinct  $t_3, t_4, t_5 \in [t_1, t_2)$ . By Lemmas B.3 and B.4,  $\text{VoteYes}_i^t \leftarrow \text{False}$  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 B.7.  $\square$

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

**Theorem B.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  $1, 2, \dots, t$ ,

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

*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

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

we know that the number of changes thus far is less than  $\min \left( \frac{\varepsilon}{8a}, \frac{\varepsilon}{4b} \right)$ , and by Lemma B.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 B.4 this means  $\text{GlobalUpdate}^t \leftarrow \text{False}$ , so by Lemma B.2 for every subgroup  $j \in [m]$  there exists user  $i \in S_j$  such that

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

Lemma B.3 then implies that

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

Since this holds for every subgroup  $j$ , we get that

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

and since  $\text{GlobalUpdate}^t \leftarrow \text{False}$ , by Lemma B.4 this means that  $\tilde{p}^t = \hat{p}^{f(t)}$  and

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

a contradiction.  $\square$

## APPENDIX C. MISSING PROOFS FROM SECTION 5.2

We start with full pseudocode for `HEAVYTHRESH` (Figure 3) as well as its subroutines (Figures 4 and 5).

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

**Lemma C.1 .**  $\mathcal{R}$  is  $\varepsilon$ -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

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

**Case 1:**  $h_1$  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.

**Case 2:** Exactly one (WLOG  $h_1$ ) hash is a nonzero vector. Then

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

Similarly,

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

Therefore

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

---

**Algorithm 4** Global Algorithm: HEAVYTHRESH

---

**Require:** number of users  $n$ , number of epochs  $T$ , minimum subgroup size  $L$ , number of subgroups  $m$ , epoch length  $\ell$ , privacy parameter  $\varepsilon$ , failure parameter  $\delta$

- 1: Initialize global estimate  $\tilde{y}^0 \leftarrow -1$
- 2: Initialize update counters  $c_1, \dots, c_n \leftarrow 0, 0, \dots, 0$
- 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  $b \leftarrow \frac{2\left(\sqrt{\frac{\ln(16wT/\delta)}{nw}} + \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}\right)}{2(\log(T)+2)\sqrt{\frac{2\ln(16wnT/\delta)}{w\ell}} - 2\sqrt{\frac{\ln(16dT/\delta)}{2wn}} - \frac{\ln(16wT/\delta)\sqrt{w}}{n^2}}$
- 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**
- 11:   GlobalUpdate $^t \leftarrow \left( \frac{1}{n} \sum_{i=1}^n a_i^t > \frac{1}{e^a+1} + \sqrt{\frac{\ln(16T/\delta)}{2n}} \right)$
- 12:   **if** GlobalUpdate $^t$  **then**
- 13:      $f(t) \leftarrow t$
- 14:     **for** each  $i \in [n]$  **do**
- 15:       User  $i$  publishes  $\tilde{z}_i^t \leftarrow \text{HEAVYEST}(i, t)$
- 16:     **end for**
- 17:     Aggregate user estimates into global estimate:  
     $\tilde{y}^t \leftarrow \frac{1}{n} \sum_{i=1}^n \tilde{z}_i^t$
- 18:   **else**
- 19:      $f(t) \leftarrow f(t-1)$
- 20:     **for** each  $i \in [n]$  **do**
- 21:       User  $i$  publishes  $\tilde{z}_i^t \leftarrow \mathcal{R}(0, b)$
- 22:     **end for**
- 23:      $\tilde{y}^t \leftarrow \tilde{y}^{t-1}$
- 24:   **end if**
- 25:   Analyst publishes  $\tilde{y}^t$
- 26:   Analyst publishes  $\text{FO}^t \leftarrow (\Phi, \tilde{y}^t)$
- 27: **end for**

---

FIGURE 3. Pseudocode for HEAVYTHRESH.

Case 3: Neither  $h_1$  nor  $h_2$  is a zero vector. Then by the logic above

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

□

---

**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:  $\text{VoteYes}_i^t \leftarrow (c_i^V < \varepsilon/4a \text{ and } 2^{\lfloor \log T \rfloor - b^*} \text{ divides } t)$
  - 5: **if**  $\text{VoteYes}_i^t$  **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:  $\mathcal{A}_{FO}$ 

---

**Require:** Frequency oracle  $(\Phi, \frac{1}{n} \sum_{i=1}^n z_i)$ , dictionary value to be estimated  $v \in [d]$ 

- 1: Output  $\hat{f}(v) = \langle \Phi e_v, \bar{z} \rangle$
- 

FIGURE 4. Pseudocode for HEAVYVOTE and  $\mathcal{A}_{FO}$ .

## APPENDIX D. MISSING PROOFS FROM SECTION 5.3

First, recall that we set

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

and

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

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

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

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

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

and for  $x \neq 0$

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

□

---

**Algorithm 7** Local Subroutine: HEAVYEST

---

**Require:** user  $i$ , epoch  $t$ 

- 1:  $\text{SendEstimate}_i^t \leftarrow \{c_i^E < \varepsilon/4b\}$
  - 2: **if**  $\text{SendEstimate}_i^t$  **then**
  - 3:    $c_i \leftarrow c_i^E + b$
  - 4:    $\tilde{z}_i^t \leftarrow \mathcal{R}(\hat{y}_i^t, b)$
  - 5: **else**
  - 6:    $\tilde{z}_i^t \leftarrow \mathcal{R}(0, b)$
  - 7: **end if**
  - 8: Output  $\tilde{z}_i^t$
- 

---

**Algorithm 8** Client Randomizer:  $\mathcal{R}$ 

---

**Require:** Hashed histogram  $h = \Phi\hat{p}_i^t$ , privacy parameter  $\varepsilon$ 

- 1: Sample  $j \in [w]$  uniformly at random
  - 2:  $c_\varepsilon \leftarrow \frac{e^b + 1}{e^b - 1}$
  - 3:  $z \leftarrow 0 \in \mathbb{R}^w$
  - 4: **if**  $h \neq 0$  **then**
  - 5:   Draw  $x \sim \text{Uni}(0, 1)$
  - 6:   **if**  $x < \frac{1}{2} + \frac{h_j \sqrt{w}}{2c_\varepsilon}$  **then**
  - 7:      $z_j \leftarrow c_\varepsilon \sqrt{w}$
  - 8:   **else**
  - 9:      $z_j \leftarrow -c_\varepsilon \sqrt{w}$
  - 10: **end if**
  - 11: **else**
  - 12:    $z_j \leftarrow c_\varepsilon \sqrt{w}$  or  $-c_\varepsilon \sqrt{w}$  u.a.r
  - 13: **end if**
  - 14: Output  $z$
- 

FIGURE 5. Pseudocode for HEAVYEST and  $\mathcal{R}$ .

Lemmas B.1 and B.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 B.3.

**Lemma D.2 .** *With probability at least  $1 - \delta/8$ , 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  $X = \frac{\sqrt{w}(\Phi\hat{p}_i^t + \frac{1}{\sqrt{w}})}{2}$  we get  $X \in [0, 1]^m$  where each index  $X_j$  has  $\mathbb{E}[X_j] = \frac{\sqrt{w}((\Phi p_{g(i)}^t)_j + \frac{1}{\sqrt{w}})}{2}$ . 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.  $\square$

This brings us to an analogue of Lemma B.4.

**Lemma D.3 .** *With probability at least  $1 - \delta/8$ , for every epoch  $t$  where every user  $i$  sets  $\text{SendEstimate}_i^t \leftarrow \text{True}$ ,*

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

*Proof.* By Lemma D.1,  $\mathbb{E}[\tilde{y}^t] = \frac{1}{n} \sum_{i=1}^n \Phi\hat{p}_i^t$ , 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  $N_s^t$  of vectors  $\hat{z}_i^t$  that are nonzero in index  $s$  by an additive Chernoff bound:

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

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  $\text{GlobalUpdate}^t \leftarrow \text{True}$

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

Similarly,

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

where the inequality holds with probability at least  $1 - \frac{\delta}{8}$  in every epoch  $t$  by the definition of  $\Phi$  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,

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

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

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

Substituting in our value

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

yields the claim.  $\square$

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 D.4 .** *With probability at least  $1 - \delta/8$ , in every epoch  $t$ , if*

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

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

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

*Proof.* The outline of our proof is similar (and in some steps identical) to that of Theorem 2.5 in (Bassily and Smith, 2015), but we provide it here for completeness. First,

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

Then by Corollary 6.4 from (Blum, Ligett and Roth, 2013), 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

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

By our assumption we can bound the first term by

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

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

We can use these lemmas to prove an analogue of Corollary B.6. First, we specify our setting-specific redefinition of *change*.

**Definition D.5 .** *We say a change  $\Delta_t$  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 D.6 .** *With probability at least  $1 - \delta$ , given adjacent changes  $\Delta_{t_1}$  and  $\Delta_{t_2}$ ,  $\text{GlobalUpdate}^t \leftarrow \text{True}$  in at most one epoch  $t \in [t_1, t_2]$ .*

*Proof.* The proof is identical to that of Lemma B.6, only replacing Lemma B.3 with Lemma D.2.  $\square$

Lemmas B.7 and B.8 apply in this setting unmodified, which finally lets us prove the following accuracy guarantee.

**Theorem D.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  $1, 2, \dots, t$ ,*

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

*Proof.* The proof is nearly identical to that of Theorem 4.3, replacing all lemmas with their heavy-hitter 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  $\min(\frac{\varepsilon}{8a}, \frac{\varepsilon}{4b})$

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

we know that the number of changes thus far is less than  $\min\left(\frac{\varepsilon}{8a}, \frac{\varepsilon}{4b}\right)$ , and by Lemma B.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 D.3 this means  $\text{GlobalUpdate}^t \leftarrow \text{False}$ , so by Lemma B.2 for every subgroup  $j \in [m]$  there exists user  $i \in S_j$  such that

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

Lemma D.2 then implies that

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

Since this holds for every subgroup  $j$ , we get that

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

and since  $\text{GlobalUpdate}^t \leftarrow \text{False}$ , by Lemma D.3 this means that  $\tilde{y}^t = \tilde{y}^{f(t)}$ , so

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

Plugging this quantity into Lemma D.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.  $\square$

