

Randomized Algorithms for Tracking Distributed Count, Frequencies, and Ranks*

Zengfeng Huang

Ke Yi

Qin Zhang

Hong Kong University of Science and Technology
{huangzf, yike}@cse.ust.hk

MADALGO, University of Aarhus
qinzhang@cs.au.dk

ABSTRACT

We show that randomization can lead to significant improvements for a few fundamental problems in distributed tracking. Our basis is the *count-tracking* problem, where there are k players, each holding a counter n_i that gets incremented over time, and the goal is to track an ε -approximation of their sum $n = \sum_i n_i$ continuously at all times, using minimum communication. While the deterministic communication complexity of the problem is $\Theta(k/\varepsilon \cdot \log N)$, where N is the final value of n when the tracking finishes, we show that with randomization, the communication cost can be reduced to $\Theta(\sqrt{k}/\varepsilon \cdot \log N)$. Our algorithm is simple and uses only $O(1)$ space at each player, while the lower bound holds even assuming each player has infinite computing power. Then, we extend our techniques to two related distributed tracking problems: *frequency-tracking* and *rank-tracking*, and obtain similar improvements over previous deterministic algorithms. Both problems are of central importance in large data monitoring and analysis, and have been extensively studied in the literature.

Categories and Subject Descriptors

F.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms

Algorithms, theory

Keywords

Distributed tracking

1. INTRODUCTION

We start with a very basic problem in distributed tracking, what we call *count-tracking*. There are k players each

*This work is supported by a DAG and an RPC grant from HKUST and a Google Faculty Research Award.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PODS'12, May 21–23, 2012, Scottsdale, Arizona, USA.

Copyright 2012 ACM 978-1-4503-1248-6/12/05 ...\$10.00.

holding a counter n_i that is initially 0. Over time, the counters get incremented and we denote by $n_i(t)$ the value of the counter n_i at time t . The goal is to track an ε -approximation of the total count $n(t) = \sum_i n_i(t)$, i.e., an $\hat{n}(t)$ such that $(1 - \varepsilon)n(t) \leq \hat{n}(t) \leq (1 + \varepsilon)n(t)$,¹ continuously at all times. There is a coordinator whose job is to maintain such an $\hat{n}(t)$, and will try to do so using minimum communication with the k players (the formal model of computation will be defined shortly).

There is a trivial solution to the count-tracking problem: Every time a counter n_i has increased by a $1 + \varepsilon$ factor, the player informs the coordinator of the change. Thus, the coordinator always has an ε -approximation of every n_i , hence an ε -approximation of their sum n . Letting N denote the final value of n , simple analysis shows that the communication cost of this algorithm is $O(k/\varepsilon \cdot \log N)$.² This algorithm was actually used in [15] for solving essentially the same problem, which also provided many practical motivations for studying this problem. Note that this algorithm is deterministic and only uses one-way communication (from the players to the coordinator), and yet it turns out this simple algorithm is already optimal for deterministic algorithms, even if two-way communication is allowed [28]. Thus the immediate questions are: What about randomized algorithms that are allowed to fail with a small probability? Is two-way communication not useful at all? In this paper, we set out to address these questions, and then move on to consider other related distributed tracking problems.

1.1 The distributed tracking model

We first give a more formal definition of the computation model that we will work with, which is essentially the same as those used in prior work on distributed tracking [2, 3, 5, 6, 8, 9, 15, 28]. There are k distributed *sites* S_1, \dots, S_k , each receiving a stream of elements over time, possibly at varying rates. Let N be the total number of elements in all k streams. We denote by $A_i(t)$ the multiset (bag) of elements received by S_i up until time t , and let $A(t) = \uplus_{i=1}^k A_i(t)$ be the combined data set, where \uplus denotes multiset addition. There is a coordinator whose job is to maintain (an approximation of) $f(A(t))$ continuously at all times, for a given function f (e.g., $f(A(t)) = |A(t)|$ for the count-tracking problem above). The coordinator has a direct two-way communication channel with each of the sites; note that broad-

¹We sometimes omit “(t)” when the context is clear.

²A more careful analysis leads to a slightly better bound of $O(k/\varepsilon \cdot \log(\varepsilon N/k))$, but we will assume that N is sufficiently large, compared to k and $1/\varepsilon$, to simplify the bounds.

casting a message costs k times the communication for a single message. The sites do not communicate with each other directly, but this is not a limitation since they can always pass messages via the coordinator. We assume that communication is instant, i.e., no element will arrive until all parties have decided not to send more messages. As in prior work, our measures of complexity will be the communication cost and the space used to process each stream. Unless otherwise specified, the unit of both measures is a *word*, and we assume that any integer less than N , as well as an element from the stream, can fit in one word.

This model was initially abstracted from many applied settings, ranging from distributed data monitoring, wireless sensor networks, to network traffic analysis, and has been extensively studied in the database community. From 2008 [8], the model has started to attract interests from the theory community as well, as it naturally combines two well-studied models: the data stream model and multi-party communication complexity. When there is only $k = 1$ site who also plays the role of the coordinator, the model degenerates to the standard streaming model; when $k \geq 2$ and our goal is to do a one-shot computation of $f(A(\infty))$, then the model degenerates to the (number-in-hand) k -party communication model. Thus, distributed tracking is more general than both models. Meanwhile, it also appears to be significantly different from either, with the above count-tracking problem being the best example. This problem is trivial in both the streaming and the communication model (even computing the exact count is trivial), whereas it becomes nontrivial in the distributed tracking model and requires new techniques, especially when randomization is allowed, as illustrated by our results in this paper.

Note that there is some work on *distributed streaming* (see e.g. [11, 16]) that adopts a model very similar to ours, but with a fundamental difference. In their model there are k streams, each of which runs a streaming algorithm on its local data. But the function f on the combined streams is computed only at the end or upon requests by the user. As one can see that the count-tracking problem is also trivial in this model. The crucial difference is that, in this model, the sites wait passively to get polled. If we want to track f continuously, we have to poll the sites all the time. Whereas in our model, the sites actively participate in the tracking protocol to make sure that f is always up-to-date.

1.2 Problem statements, previous and new results

In this paper, we first study the count-tracking problem. Then we extend our approach to two related, more general problems: *frequency-tracking* and *rank-tracking*. Both problems are of central importance in large data monitoring and analysis, and have been extensively studied in the literature. In all the communication upper bounds, we will assume $k \leq 1/\varepsilon^2$; otherwise all of them will carry an extra additive $O(k \log N)$ term. All our results are summarized in Table 1; below we discuss each of them separately.

As mentioned earlier, the deterministic communication complexity for the count-tracking problem has been settled at $\Theta(k/\varepsilon \cdot \log N)$ [28]³, with or without two-way communication. In this paper, we show that with randomization and

two-way communication, this is reduced to $\Theta(\sqrt{k}/\varepsilon \cdot \log N)$. We first in Section 2.1 present a randomized algorithm with this communication cost that, at *any one* given time instance, maintains an ε -approximation of the current n with a constant probability. The algorithm is very simple and uses $O(1)$ space at each site. It is easy to make the algorithm correct for *all* time instances and boost the probability to $1 - \delta$: Since we can use the same approximate value \hat{n} of n until n grows by a $1 + \varepsilon$ factor, it suffices to make the algorithm correct for $O(\log_{1+\varepsilon} N) = O(1/\varepsilon \cdot \log N)$ time instances. Then running $O(\log(\frac{\log N}{\delta\varepsilon}))$ independent copies of the algorithm and taking the median will achieve the goal of tracking n continuously at all times, with probability at least $1 - \delta$. The $\Omega(\sqrt{k}/\varepsilon \cdot \log N)$ lower bound (Section 2.2) actually holds on the number of messages that have to be exchanged, regardless of the message size, and holds even assuming the sites have unlimited space and computing power. That randomization is necessary to achieve this \sqrt{k} -factor improvement follows from the previous deterministic lower bound [28]; here in Section 2.2 we give a proof that two-way communication is also required. More precisely, we show that any randomized algorithm with one-way communication has to use $\Omega(k/\varepsilon \cdot \log N)$ communication, i.e., the same as that for deterministic algorithms.

In the *frequency-tracking* (a.k.a. *heavy hitters tracking*) problem, $A(t)$ is a multiset of cardinality $n(t)$ at time t . Let $f_j(t)$ be the frequency of element j in $A(t)$. The goal is to maintain a data structure from which $f_j(t)$, for any given j , can be estimated with absolute error at most $\varepsilon n(t)$, with probability at least 0.9 (say). Note that this problem degenerates to count-tracking when there is only one element. It is reasonable to ask for an error in terms of $n(t)$: if the error were $\varepsilon f_j(t)$, then every element would have to be reported if they were all distinct. In fact, this error requirement is the widely accepted definition for the heavy hitters problem, which has been extensively studied in the streaming literature [7]. Several algorithms with the optimal $O(1/\varepsilon)$ space exist [17–19]. In the distributed tracking model, we previously [28] gave a deterministic algorithm with $O(k/\varepsilon \cdot \log N)$ communication, which is the best possible for deterministic algorithms. In this paper, by generalizing our count-tracking algorithm, we reduce the cost to $O(\sqrt{k}/\varepsilon \cdot \log N)$, with randomization (Section 3). Since this problem is more general than count-tracking, by the count-tracking lower bound, this is also optimal. Our algorithm uses $O(1/(\varepsilon\sqrt{k}))$ space to process the stream at each site, which is actually smaller than the $\Omega(1/\varepsilon)$ space lower bound for this problem in the streaming model. This should not come as a surprise: Due to the fact that the site is allowed to communicate to the coordinator *during* the streaming process, the streaming lower bounds do not apply in our model. To this end, we prove a new space lower bound of $\Omega(1/(\varepsilon\sqrt{k}))$ bits for our model, showing that our algorithm also uses near-optimal space. This space lower bound is conditioned upon the requirement that the communication cost should be $O(\sqrt{k}/\varepsilon \cdot \log N)$ bits. Note that it is not possible to prove a space lower bound unconditional of communication: A site can send every element to the coordinator and thus only needs $O(1)$ space. In fact, what we prove is a space-communication trade-off; please see Section 3.2 for the precise statement.

For the *rank-tracking* problem, it will be convenient to assume that the elements are drawn from a totally ordered universe and $A(t)$ contains no duplicates. The *rank* of an

³The lower bound in [28] was stated for the heavy hitters tracking problem, but essentially the same proof works for count-tracking.

		space (per site)	communication
count-tracking	trivial	$O(1)$	$\Theta(k/\varepsilon \cdot \log N)$
	new	$O(1)$	$O(\sqrt{k}/\varepsilon \cdot \log N)$ $\Omega(\sqrt{k}/\varepsilon \cdot \log N)$ messages
frequency-tracking	[28]	$O(1/\varepsilon)$	$\Theta(k/\varepsilon \cdot \log N)$
	new	$O(1/(\varepsilon\sqrt{k}))$ $\Omega(1/(\varepsilon\sqrt{k}))$ bits*	$O(\sqrt{k}/\varepsilon \cdot \log N)$ $\Omega(\sqrt{k}/\varepsilon \cdot \log N)$ messages
rank-tracking	[28]	$O(1/\varepsilon \cdot \log n)$	$O(k/\varepsilon \cdot \log N \log^2(1/\varepsilon))$
	new	$O\left(1/(\varepsilon\sqrt{k}) \cdot \log^{1.5} \frac{1}{\varepsilon} \log^{0.5} \frac{1}{\varepsilon\sqrt{k}}\right)$ $\Omega(1/(\varepsilon\sqrt{k}))$ bits*	$O\left(\sqrt{k}/\varepsilon \cdot \log N \log^{1.5} \frac{1}{\varepsilon\sqrt{k}}\right)$ $\Omega(\sqrt{k}/\varepsilon \cdot \log N)$ messages
sampling	[9]	$O(1)$	$O(1/\varepsilon^2 \cdot \log N)$

Table 1: Space and communication costs of previous and new algorithms. We assume $k \leq 1/\varepsilon^2$. All upper bounds are in terms of words. *This is conditioned upon the communication cost being $O(\sqrt{k}/\varepsilon \cdot \log N)$ bits.

element x in $A(t)$ (x need not be in $A(t)$) is the number of elements in $A(t)$ smaller than x , and our goal is to compute a data structure from which the rank of any given x can be estimated with error at most $\varepsilon n(t)$, with constant probability. Note that a rank-tracking algorithm also solves the frequency-tracking problem (but not vice versa), by turning each element x into a pair (x, y) to break all ties (by comparing the pairs lexicographically) and maintaining such a rank-tracking data structure. When the frequency of x is desired, we ask for the ranks of $(x, 0)$ and (x, ∞) and take the difference. We previously [28] gave a deterministic algorithm for the rank-tracking problem with communication $O(k/\varepsilon \cdot \log N \log^2(1/\varepsilon))$. In this paper, we show in Section 4 how randomization can bring this down to $O(\sqrt{k}/\varepsilon \cdot \log N \log^{1.5}(1/\varepsilon\sqrt{k}))$, which is again optimal ignoring polylog($1/\varepsilon, k$) factors. Since rank-tracking is more general than frequency-tracking, the previous lower bounds also hold here. Our algorithm uses space that is also close to the $\Omega(1/(\varepsilon\sqrt{k}))$ lower bound.

Since we are talking about randomized algorithms with a constant success probability, we should also compare with random sampling. It is well known [24] that this probabilistic guarantee can be achieved for all the problems above by taking a random sample of size $O(1/\varepsilon^2)$. A random sample can be maintained continuously over distributed streams [9, 23], solving these distributed tracking problems, with a communication cost of $O(1/\varepsilon^2 \cdot \log N)$. This is worse than our algorithms when $k = o(1/\varepsilon^2)$. As noted earlier, all the upper bounds we have mentioned above have a hidden additive $O(k \log N)$ term. Thus when $k = \Omega(1/\varepsilon^2)$, all of them boil down to $O(k \log N)$,⁴ so our results are more interesting for the $k \leq 1/\varepsilon^2$ case, which we will assume in all the upper bounds throughout the paper. In the lower bound statements, however, we do distinguish between the two cases. The lower bounds in Table 1 assume $k \leq 1/\varepsilon^2$, and they all match the upper bounds (except for the rank-tracking problem); for the $k = \Omega(1/\varepsilon^2)$ case, the lower bound is $\Omega(k)$ (Theorem 2.3), which leaves a gap of $\Theta(\log N)$ from the upper bound.

The idea behind all our algorithms is very simple. Instead of deterministic algorithms, we use randomized algorithms that produce unbiased estimators for n_i , the frequencies, and ranks with variance $(\varepsilon n)^2/k$, leading to an overall variance

⁴The bound of the random sampling algorithm [9, 23] is actually slightly better, which is $O(k \log N / \log(k\varepsilon^2))$.

of $(\varepsilon n)^2$, which is sufficient to produce an estimate within error εn with constant probability. This means we can afford an error of $\varepsilon n/\sqrt{k}$ from each site, as opposed to $\varepsilon n/k$ for deterministic algorithms. This is essentially where we obtain the \sqrt{k} -factor improvement by randomization. Our algorithms are simple and extremely lightweight, in particular the count-tracking and frequency-tracking algorithms, thus can be easily implemented in power-limited distributed systems like wireless sensor networks.

1.3 Other related work

As distributed tracking is closely related to the streaming and the k -party communication model, it could be enlightening to compare with the known results of the above problems in these models. As mentioned earlier, the count-tracking problem is trivial in both models, requiring $O(1)$ space in the streaming model and $O(k)$ communication in the k -party communication model.

Both the frequency-tracking and rank-tracking problems have been extensively studied in the streaming model with a long history. The former was first resolved by the MG algorithm [19] with the optimal space $O(1/\varepsilon)$, though several other algorithms with the same space bound have been proposed later on [17, 18]. The rank problem is also one of the earliest problems studied in the streaming model [20]. The best deterministic algorithm to date is the one by Greenwald and Khanna [12]. It uses $O(1/\varepsilon \cdot \log n)$ working space to maintain a structure of size $O(1/\varepsilon)$, from which any rank can be estimated with error εn . Note that the rank problem is often studied as the *quantiles* problem in the literature. Recall that for any $0 \leq \phi \leq 1$, the ϕ -quantile of D is the element in $A(t)$ that ranks at $\lfloor \phi n \rfloor$, while an ε -approximate ϕ -quantile is any element that ranks between $(\phi - \varepsilon)n$ and $(\phi + \varepsilon)n$. Clearly, if we have the data structure for one problem, we can do a binary search to solve the other. Thus the two problems are equivalent, for deterministic algorithms. For algorithms with probabilistic guarantees, we need all $O(\log(1/\varepsilon))$ decisions in the binary search to succeed, which requires the failure probability to be lowered by an $O(\log(1/\varepsilon))$ factor. By running $O(\log \log(1/\varepsilon))$ independent copies of the algorithm, this is not a problem. So the two problems differ by at most a factor of $O(\log \log(1/\varepsilon))$.

The existing streaming algorithms for the frequency and rank problems can be used to solve the one-shot version of the problem in the k -party communication model easily. More precisely, we use a streaming algorithm to sum-

marize the data set at each site with a structure of size $O(1/\varepsilon)$, and then send these summary structures to the coordinator, resulting in a communication cost of $O(k/\varepsilon)$. Recently, we designed randomized algorithms for these two problems with $O(\sqrt{k}/\varepsilon)$ communication [13, 14], which have just been shown to be near-optimal [26]. Here we have extended the one-shot algorithms of [13, 14] to the continuous tracking setting. The results have demonstrated that, the seemingly more challenging tracking problem, which requires us to solve the one-shot problem continuously at all times, is only harder by an $\Theta(\log N)$ factor than the one-shot version (except for the count-tracking problem, which is much harder than its one-shot version).

Finally, we should mention that all these distributed tracking problems have been studied in the database community previously, but mostly using heuristics. Keralapura et al. [15] approached the count-tracking problem using prediction models, which do not work under adversarial inputs. Babcock and Olston [3] studied the top- k tracking problem, a variant of the frequency (heavy hitters) tracking problem, but did not offer a theoretical analysis. The rank-tracking problem was first studied by Cormode et al. [6]; their algorithm has a communication cost of $O(k/\varepsilon^2 \cdot \log N)$ under certain inputs.

2. TRACKING DISTRIBUTED COUNT

2.1 The algorithm

The algorithm with a fixed p .

Let p be a parameter to be determined later. For now we will assume that p is fixed. The algorithm is very simple: Whenever site S_i receives an element (hence n_i gets incremented by one), it sends the latest value of n_i to the coordinator with probability p . Let \bar{n}_i be the last updated value of n_i received by the coordinator. We first estimate each n_i by

$$\hat{n}_i = \begin{cases} \bar{n}_i - 1 + 1/p, & \text{if } \bar{n}_i \text{ exists;} \\ 0, & \text{else.} \end{cases} \quad (1)$$

Then we estimate n as $\hat{n} = \sum_i \hat{n}_i$.

Analysis.

As mentioned in the introduction, our analysis will hold for any given one time instance. It is also important to note that this given time instance shall not depend on the randomization internal to the algorithm.

We show that each \hat{n}_i is an unbiased estimator of n_i with variance at most $1/p^2$. This is very intuitive, since $n_i - \bar{n}_i$ is the number of failed trials until the site decides to send an update to the coordinator, when we look backward from the current time instance. This follows a geometric distribution with parameter p , but not quite, as it is bounded by n_i . This is why we need to separate the two cases in (1). The following calculation appeared in [13]; we include it here for completeness.

LEMMA 2.1. $\mathbf{E}[\hat{n}_i] = n_i; \mathbf{Var}[\hat{n}_i] \leq 1/p^2$.

Proof. Define the random variable

$$X = \begin{cases} n_i - \bar{n}_i + 1, & \text{if } \bar{n}_i \text{ exists;} \\ n_i + 1/p, & \text{else.} \end{cases}$$

Now we can rewrite \hat{n}_i as $\hat{n}_i = n_i - X + 1/p$. Thus it suffices to show that $\mathbf{E}[X] = 1/p$ and $\mathbf{Var}[X] \leq 1/p^2$. Letting $t = n_i - \bar{n}_i + 1$, we have

$$\begin{aligned} \mathbf{E}[X] &= \sum_{t=1}^{n_i} (t(1-p)^{t-1}p) + (n_i + 1/p)(1-p)^{n_i} = \frac{1}{p}. \\ \mathbf{Var}[X] &= \sum_{t=1}^{n_i} ((t-1/p)^2(1-p)^{t-1}p) + (n_i + 1/p - 1/p)^2(1-p)^{n_i} \\ &= \frac{(1-p)(1-(1-p)^{n_i})}{p^2} \leq \frac{1}{p^2}. \end{aligned}$$

□

By Lemma 2.1, we know that \hat{n} is an unbiased estimator of n with variance $\leq k/p^2$. Thus, if $p = \sqrt{k}/\varepsilon n$, the variance of \hat{n} will be $(\varepsilon n)^2$, which means that \hat{n} has error at most $2\varepsilon n$ with probability at least $3/4$, by Chebyshev inequality. Rescaling ε and p by a constant will reduce the error to εn and improves the success probability to 0.9, as desired.

Previously, we used similar ideas to solve the *one-shot* quantile problem over distributed data [13]. Here, we essentially treat the numerical values of the items in the quantile problem as the timestamps in the count-tracking problem, and simulate the one-shot sampling algorithm of [13] in the continuous setting.

Dealing with a decreasing p .

It is not possible and necessary to set p exactly to $\sqrt{k}/\varepsilon n$. From the analysis above, it should be clear that keeping $p = \Theta(\sqrt{k}/\varepsilon n)$ will suffice. To do so, we first track n within a constant factor. This can be done efficiently as follows. Each site S_i keeps track of its own counter n_i . Whenever n_i doubles, it sends an update to the coordinator. The coordinator sets $n' = \sum_{i=1}^k n'_i$, where n'_i is the last update of n_i . When n' doubles (more precisely, when n' changes by a factor between 2 and 4), the coordinator broadcasts n' to all the sites. Let \bar{n} be the last broadcast value of n' . It is clear that \bar{n} is always a constant-factor approximation of n . The communication cost is $O(k \log N)$, since each site sends $O(\log N)$ updates to the coordinator and the coordinator broadcasts $O(\log N)$ times, each of which costs k messages. These broadcasts divide the whole tracking period into $O(\log N)$ rounds, and within each round, n stays within a constant factor of \bar{n} , the broadcast value at the beginning of the round.

Now, when $\bar{n} \leq \sqrt{k}/\varepsilon$, all the sites set $p = 1$. This causes all the first $O(\sqrt{k}/\varepsilon)$ elements to be sent to the coordinator. When $\bar{n} > \sqrt{k}/\varepsilon$, the sites use $p = 1/\lceil \varepsilon \bar{n} / \sqrt{k} \rceil_2$, where $\lfloor x \rfloor_2$ denotes the largest power of 2 smaller than x . Since \bar{n} is monotonically increasing, p gets halved over the rounds. At the beginning of a round, if the new p is half⁵ of that in the previous round, each site S_i adjusts its \bar{n}_i appropriately, as follows. First with probability $1/2$, the site decides if \bar{n}_i remains the same. If so, nothing changes; otherwise, it repeatedly flips a coin with probability $1/p$ (with the new p). Every failed coin flip decrements \bar{n}_i by one. It does so until a successful coin flip, or $\bar{n}_i = 0$. Finally, the site informs the coordinator of the new value of \bar{n}_i ; if $\bar{n}_i = 0$, the coordinator will treat it as if \bar{n}_i does not exist. It should be clear that

⁵To be more precise, the new p might also be a quarter of the previous p , but it can be handled similarly.

after this adjustment, the whole system looks as if it had always been running with the new p .

It is easy to see that the communication cost in each round is $O(k + pn) = O(k + \sqrt{k}/\varepsilon) = O(\sqrt{k}/\varepsilon)$, thus the total cost is $O(\sqrt{k}/\varepsilon \cdot \log N)$.

THEOREM 2.1. *There is an algorithm for the count-tracking problem that, at any time, estimates $n = \sum_i n_i$ within error εn with probability at least 0.9. It uses $O(1)$ space at each site and $O(\sqrt{k}/\varepsilon \cdot \log N)$ total communication.*

2.2 The lower bound

Before proving the lower bounds, we first state our lower bound model formally, in the context of the count-tracking problem. The N elements arrive at the k sites in an online fashion at arbitrary time instances. We do not allow spontaneous communication. More precisely, it means that a site is allowed to send out a message only if it has just received an element or a message from the coordinator. Likewise, the coordinator is allowed to send out messages only if it has just received messages from one or more sites. When a site S_j is allowed to send out a message, it decides whether it will indeed send a message, and the content of the message if so, based only on its local counter n_j and the message history between S_j and the coordinator, possibly using some random source. We assume that the site does not look at the current clock. We argue that the clock conveys no information since the elements arrive at arbitrary and unpredictable time instances. (If the elements arrive in a predictable fashion, say, one per time step, the problem can be solved without communication at all.) Similarly, when the coordinator is allowed to send out messages, it makes the decision on where and what to send based only on its message history and some random source. We will lower bound the communication cost only by the number of messages, regardless of the message size.

2.2.1 One-way communication lower bound

In this section we show that two-way communication is necessary to achieve the upper bound in Theorem 2.1, by proving the following lower bound. Remember that we assume N is sufficiently larger than k and $1/\varepsilon$.

THEOREM 2.2. *If only the sites can send messages to the coordinator but not vice versa, then any randomized algorithm for the count-tracking problem that, at any time, estimates n within error εn with probability at least 0.9 must send $\Omega(k/\varepsilon \cdot \log N)$ messages.*

Proof. We first define the hard input distribution μ .

- (a) With probability $1/2$, all elements arrive at one site that is uniformly picked at random.
- (b) Otherwise, the N elements arrive at the k sites in a round-robin fashion, each site receiving N/k elements in the end.

By Yao's Minimax principle [27], we only need to argue that any deterministic algorithm with success probability at least 0.8 under μ has expected cost $\Omega(k/\varepsilon \cdot \log N)$.

Note that when only one-way communication is allowed, a site decides whether to send messages to the coordinator only based on its local counter n_j . Thus the communication

pattern can be essentially described as follows. Each site S_j has a series of thresholds t_j^1, t_j^2, \dots such that when $n_j = t_j^i$, the site sends the i -th message to the coordinator. These thresholds should be fixed at the beginning.

We lower bound the communication cost by rounds. Let W_i be the number of elements that have arrived up until round i . We divide the rounds by setting $W_1 = k/\varepsilon$, and $W_{i+1} = \lceil (1+\varepsilon)W_i \rceil$ for $i \geq 1$. Thus there are $1/\varepsilon \cdot \log(\varepsilon N/k)$ rounds, which is $\Omega(1/\varepsilon \cdot \log N)$ for sufficiently large N .

At the beginning of round $i+1$, suppose that S_1, S_2, \dots, S_k have already sent $z_1^i, z_2^i, \dots, z_k^i$ messages to the coordinator, respectively. Let $t_{\max}^{i+1} = (1+\varepsilon) \cdot \max\{t_j^{z_j^i} \mid j = 1, 2, \dots, k\}$. We first observe that there must be at least $k/2$ sites with their next threshold $t_j^{z_j^i+1} \leq t_{\max}^{i+1}$. Otherwise, suppose there are less than $k/2$ sites with such next thresholds, then with probability at least $1/4$ case (a) happens and the random site S_j chosen to receive all elements has $t_j^{z_j^i+1} > t_{\max}^{i+1} \geq (1+\varepsilon)t_j^{z_j^i}$. Thus, with probability at least $1/4$ the algorithm fails when the t_{\max}^{i+1} -th element arrives, contradicting the success guarantee.

On the other hand, with probability $1/2$ case (b) happens.

In this case all $t_j^{z_j^i}$ ($j = 1, 2, \dots, k$) are no more than W_i/k , since in case (b), elements arrive at all k sites in turn. In the next εW_i elements, each site S_j receives $\varepsilon W_i/k$ elements. If the site S_j has $t_j^{z_j^i+1} \leq t_{\max}^{i+1}$, then it must send a message in this round, since $W_i/k + \varepsilon W_i/k \geq t_{\max}^{i+1} \geq t_j^{z_j^i+1}$, that is, its $(z_j^i + 1)$ -th threshold is triggered. As argued, there are $\geq k/2$ sites with $t_j^{z_j^i+1} \leq t_{\max}^{i+1}$, so the communication cost in this round is at least $k/2$.

Summing up all rounds, the total communication is at least $\Omega(k/\varepsilon \cdot \log N)$. \square

2.2.2 Two-way communication lower bound

Below we prove two randomized lower bounds when two-way communication is allowed. The first one justifies the assumption $k \leq 1/\varepsilon^2$, since otherwise, random sampling will be near-optimal.

THEOREM 2.3. *Any randomized algorithm for the count-tracking problem that, at any time, estimates n within error $0.1n$ with probability at least 0.9 must exchange $\Omega(k)$ messages.*

Proof. The hard input distribution is the same as that in the proof of Theorem 2.2. To prove this lower bound we are only interested in the number of sites that communicate with the coordinator at least once. Before any element arrives, we can still assume that each site keeps a triggering threshold. The thresholds of S_j shall remain the same unless it communicates with the coordinator at least once. We argue that there must be at least $k/2$ sites whose triggering threshold is no more than 1, since otherwise if case (a) happens and the randomly chosen site is one with a triggering threshold larger than 1, the algorithm will fail, which would happen with probability at least $1/4$. On the other hand, if case (b) happens, then all the sites with threshold 1 will have to communicate with the coordinator at least once: either their thresholds are triggered by the round-robin arrival of elements, or they receive a message from the coordinator,

which can possibly change their threshold. \square

Finally, we show that the upper bound in Theorem 2.1 is tight. We first introduce the following primitive problem.

DEFINITION 2.1 (1-BIT). *Let s be either $k/2 + \sqrt{k}$ or $k/2 - \sqrt{k}$, each with probability $1/2$. From the k sites, a subset of s sites picked uniformly at random each have bit 1, while the other $k - s$ sites have bit 0. The goal of the communication problem is for the coordinator to find out the value of s with probability at least 0.8.*

We will show the following lower bound for the 1-BIT problem.

LEMMA 2.2. *Any deterministic algorithm that solves 1-bit has distributional communication complexity $\Omega(k)$.*

Lemma 2.2 immediately implies the following theorem:

THEOREM 2.4. *Any randomized algorithm for the count-tracking problem that, at any time, estimates n within error εn with probability at least 0.9 must exchange $\Omega(\sqrt{k}/\varepsilon \cdot \log N)$ messages, when $k < 1/\varepsilon^2$.*

Proof. We will again fix a hard input distribution first and then focus on the distributional communication complexity of deterministic algorithms with success probability at most 0.8. Let $[m] = \{0, 1, \dots, m - 1\}$. The adversarial input consists of $\ell = \log \frac{\varepsilon N}{k} = \Omega(\log N)$ rounds. We further divide each round $i \in [\ell]$ into $r = 1/(2\varepsilon\sqrt{k})$ subrounds.

The input at round $i \in [\ell]$ is constructed as follows, at each subround $j \in [r]$, we first choose s to be $k/2 + \sqrt{k}$ or $k/2 - \sqrt{k}$ with equal probability. Then we choose s sites out of the k sites uniformly at random and send 2^i elements to each of them (the order does not matter).

It is easy to see that at the end of in each subround in round i , the total number of items is no more than $\tau_i = \sqrt{k}/\varepsilon \cdot 2^i$. Thus after $s \cdot 2^i$ elements have arrived in a subround, the algorithm has to correctly identify the value of s with probability at least 0.8, since otherwise with probability at least 0.2 the estimation of the algorithm will deviate from the true value by at least $\sqrt{k} \cdot 2^i > \varepsilon \tau_i$, violating the success guarantee of the algorithm. This is exactly the 1-bit problem defined above. By Lemma 2.2, the communication cost of each subround is $\Omega(k)$. Summing over all r subrounds and then all ℓ rounds, we have that the total communication is at least $\ell \cdot r \cdot \Omega(k) \geq \Omega(\sqrt{k}/\varepsilon \cdot \log N)$. \square

Now we prove Lemma 2.2.

Proof. (of Lemma 2.2) First of all, observe that whenever the coordinator communicates with a site, the site can send its whole input (i.e., its only bit) to the coordinator. After that, the coordinator knows all the information about that site and does not need to communicate with it further. Therefore all that we need to investigate is the number of sites the coordinator needs to communicate with.

There can be two types of actions in the protocol.

- (a) A site initiates a communication with the coordinator based on the bit it has.
- (b) The coordinator, based on all the information it has gathered so far, asks some site to send its bit.

Note that if a type (b) communication takes place before a type (a) communication, we can always swap the two, since this only gives the coordinator more information at an earlier stage. Thus we can assume that all the type (a) communications happen before type (b) ones.

In the first phase where all the type (a) communications happen, let x be the number of sites that send bit 0 to the coordinator, and y be the number of sites that send bit 1 to the coordinator. If $\mathbb{E}[x + y] = \Omega(k)$, then we are done. So let us assume that $\mathbb{E}[x + y] = o(k)$. By Markov inequality we have that, with probability at least 0.9, $x + y = o(k)$. After the first phase, the problem becomes that there are $s' = s - y = s - o(k)$ sites having bit 1, out of a total $k' = k - x - y = k - o(k)$ sites. The coordinator needs to figure out the exact value of s' with probability at least $0.8 - (1 - 0.9) = 0.7$.

In the second phase where all type (b) communication happens, from the coordinator's perspective, all the remaining sites are still symmetric (by the random input we choose), therefore the best it can do is to probe an arbitrary site among those that it has not communicated with. This is still true even after the coordinator has probed some of the remaining sites. Therefore, the problem boils down to the following: The coordinator picks z sites out of the remaining k' sites to communicate and then decides the value of s' with success probability at least 0.7. We call this problem the *sampling* problem. We can show that to achieve the success guarantee, z should be at least $\Omega(k)$. This result is perhaps folklore; proofs to more general versions of this problem can be found in [4] (Chapter 4), and also [21, 25]. We include a simpler proof in the appendix for completeness. With this we conclude the proof of Lemma 2.2. \square

3. TRACKING DISTRIBUTED FREQUENCIES

In the frequency-tracking problem, A (we omit “(t)” when the context is clear) is a multiset and the goal is to track the frequency of any item j within error εn . Let f_{ij} denote the local frequency of element j in A_i , and let $f_j = \sum_{i=1}^k f_{ij}$.

3.1 The algorithm

The algorithm with a fixed p .

As in Section 2.1 we first describe the algorithm with a fixed parameter p . If each site tracks the local frequencies f_{ij} exactly, we can essentially use the count-tracking algorithm to track the f_j 's. To achieve small space, we make use of the following algorithm due to Manku and Motwani [17] at each site S_i : The site maintains a list L_i of counters. When an element j arrives at S_i , the site first checks if there is a counter c_{ij} for j in L_i . If yes, it increases c_{ij} by 1. Otherwise, the site samples this element with probability p . If it is sampled, the site inserts a counter c_{ij} , initialized to 1, into L_i . It is easy to see that the expected size of L_i is $O(pn_i)$.

Next, we follow a similar strategy as in the count-tracking algorithm: The site reports the counter c_{ij} to the coordinator when it is first added to the counter list with an initial value of 1. Afterward, for every j that is arriving, the site always increments c_{ij} as before, but only sends the updated counter to the coordinator with probability p . We use \bar{c}_{ij} to denote the last updated value of c_{ij} .

The tricky part is how the coordinator estimates f_{ij} , hence f_j . Fix any time instance. The difference between f_{ij} and \hat{c}_{ij} comes from two sources: one is the number of j 's missed before a copy is sampled, and the other is the number of j 's that arrive after the last update of c_{ij} . It is easy to see that both errors follow the same distribution as $n_i - \bar{n}_i$ in the count-tracking algorithm. Thus it is tempting to modify (1) as

$$\hat{f}_{ij} = \begin{cases} \bar{c}_{ij} - 2 + 2/p, & \text{if } \bar{c}_{ij} \text{ exists;} \\ 0, & \text{else.} \end{cases} \quad (2)$$

However, this estimator is biased and its bias might be as large as $\Theta(\varepsilon n/\sqrt{k})$. Summing over k streams, this would exceed our error guarantee. To see this, consider the f_{ij} copies of j . Effectively, the site samples every copy with probability p , while $\bar{c}_{ij} - 2$ is exactly the number of copies between the first and the last sampled copy (excluding both). We define X_1 as before

$$X_1 = \begin{cases} t_1, & \text{if the } t_1\text{th copy is the first one sampled;} \\ f_{ij} + 1/p, & \text{if none is sampled.} \end{cases}$$

We define X_2 in exactly the same way, except that we examine these f_{ij} copies backward:

$$X_2 = \begin{cases} t_2, & \text{if the } t_2\text{th copy is the first one sampled} \\ & \text{in the reverse order;} \\ f_{ij} + 1/p, & \text{if none is sampled.} \end{cases}$$

It is clear that X_1 and X_2 have the same distribution with $\mathbb{E}[X_1] = \mathbb{E}[X_2] = 1/p$ (by Lemma 2.1), so $\hat{f}_{ij} = f_{ij} - (X_1 + X_2) + 2/p$ is unbiased. Since $\bar{c}_{ij} - 2 = f_{ij} - t_1 - t_2$, the correct unbiased estimator should be

$$\hat{f}_{ij} = \begin{cases} \bar{c}_{ij} - 2 + 2/p, & \text{if } \bar{c}_{ij} \text{ exists;} \\ -f_{ij}, & \text{else.} \end{cases} \quad (3)$$

Compared with the previous wrong estimator (2), the main difference is how the estimation is done when no copy of j is sampled. When $f_{ij} = \Theta(\varepsilon n/\sqrt{k})$ and $p = \Theta(1/f_{ij})$, this happens with constant probability, which would result in a bias of $\Theta(f_{ij}) = \Theta(\varepsilon n/\sqrt{k})$.

However, the correct estimator (3) depends on f_{ij} , the quantity we want to estimate in the first place. The workaround is to use another unbiased estimator for f_{ij} when \bar{c}_{ij} is not yet available. It turns out that we can just use simple random sampling: The site samples every element with probability p (this is independent of the sampling process that maintains the list L_i), and sends the sampled elements to the coordinator. Let d_{ij} be the number of sampled copies of j received by the coordinator from site i , the final estimator for f_{ij} is

$$\hat{f}'_{ij} = \begin{cases} \bar{c}_{ij} - 2 + 2/p, & \text{if } \bar{c}_{ij} \text{ exists;} \\ -d_{ij}/p, & \text{else.} \end{cases} \quad (4)$$

Since d_{ij} is independent of \bar{c}_{ij} , the estimator is still unbiased. Below we analyze its variance.

Analysis.

Intuitively, the variance is not affected by using the simple random sampling estimator d_{ij}/p , because it is only used when \bar{c}_{ij} is not available, which means that f_{ij} is likely to be small, and when f_{ij} is small, d_{ij}/p actually has a small variance. When f_{ij} is large, d_{ij}/p has a large variance, but we will use it only with small probability. Below we give a formal proof.

LEMMA 3.1. $\mathbb{E}[\hat{f}'_{ij}] = f_{ij}$; $\text{Var}[\hat{f}'_{ij}] = O(1/p^2)$.

Proof. We first analyze the estimator \hat{f}_{ij} of (3). That $\mathbb{E}[\hat{f}_{ij}] = f_{ij}$ follows from the discussion above. Its variance is $\text{Var}[\hat{f}_{ij}] = \text{Var}[X_1 + X_2]$. Note that X_1 and X_2 are not independent, but they both have expectation $1/p$ and variance $\leq 1/p^2$. We first rewrite

$$\begin{aligned} \text{Var}[X_1 + X_2] &= \mathbb{E}[X_1^2 + X_2^2 + 2X_1X_2] - \mathbb{E}[X_1 + X_2]^2 \\ &= \text{Var}[X_1] + \mathbb{E}[X_1^2] + \text{Var}[X_2] + \mathbb{E}[X_2^2] \\ &\quad + 2\mathbb{E}[X_1X_2] - (\mathbb{E}[X_1] + \mathbb{E}[X_2])^2 \\ &\leq 4/p^2 + 2\mathbb{E}[X_1X_2] - 4/p^2 \leq 2\mathbb{E}[X_1X_2]. \end{aligned}$$

Let \mathcal{E}_t be the event that the t th copy of j is the first being sampled. We have

$$\begin{aligned} \mathbb{E}[X_1X_2] &= \sum_{t=1}^{f_{ij}} (1-p)^{t-1} p t \mathbb{E}[X_2 | \mathcal{E}_t] + (1-p)^{f_{ij}} (f_{ij} + 1/p)^2 \\ &= \sum_{t=1}^{f_{ij}} (1-p)^{t-1} p t \left((1-p)^{f_{ij}-t} (f_{ij} - t + 1) + \sum_{l=1}^{f_{ij}-t} (1-p)^{l-1} p l \right) \\ &\quad + (1-p)^{f_{ij}} (f_{ij} + 1/p)^2 \\ &\leq \frac{1}{p^2} + (1-p)^{f_{ij}} f_{ij}^2 + \frac{(1-p)^{f_{ij}} f_{ij}}{p}. \end{aligned}$$

Let $c = f_{ij}p$. If $c \leq 2$, $f_{ij} \leq 2/p$, and the variance is $O(1/p^2)$. Otherwise

$$\mathbb{E}[X_1X_2] \leq \frac{1}{p^2} + \frac{c^2}{p^2 e^c} + \frac{c}{p^2 e^c} = O(1/p^2),$$

since $c^2 \leq e^c$ when $c > 2$.

Next we analyze the final estimator \hat{f}'_{ij} of (4). First, d_{ij} is the sum of f_{ij} Bernoulli random variables with probability p , so $\mathbb{E}[d_{ij}/p] = f_{ij}$ and $\text{Var}[d_{ij}/p] \leq f_{ij}p/p^2 = f_{ij}/p$. Let \mathcal{E}_* be the event that \hat{c}_{ij} is available, i.e., at least one copy of j is sampled, and $\mathcal{E}_0 = \bar{\mathcal{E}}_*$, then

$$\begin{aligned} \mathbb{E}[\hat{f}'_{ij}] &= \mathbb{E}[\hat{f}_{ij} | \mathcal{E}_*] \Pr[\mathcal{E}_*] + \mathbb{E}[-d_{ij}/p | \mathcal{E}_0] \Pr[\mathcal{E}_0] \\ &= \mathbb{E}[\hat{f}_{ij} | \mathcal{E}_*] \Pr[\mathcal{E}_*] + (-f_{ij}) \Pr[\mathcal{E}_0] \\ &= \mathbb{E}[\hat{f}_{ij}] = f_{ij}. \end{aligned}$$

The variance is

$$\begin{aligned} \text{Var}[\hat{f}'_{ij}] &= \mathbb{E}[\hat{f}'_{ij}{}^2] - \mathbb{E}[\hat{f}'_{ij}]^2 \\ &= \mathbb{E}[\hat{f}_{ij}^2 | \mathcal{E}_*] \Pr[\mathcal{E}_*] + \mathbb{E}[(d_{ij}/p)^2 | \mathcal{E}_0] \Pr[\mathcal{E}_0] - f_{ij}^2 \\ &= \mathbb{E}[\hat{f}_{ij}^2 | \mathcal{E}_*] \Pr[\mathcal{E}_*] - f_{ij}^2 + \mathbb{E}[(d_{ij}/p)^2] \Pr[\mathcal{E}_0] \\ &= \mathbb{E}[\hat{f}_{ij}^2 | \mathcal{E}_*] \Pr[\mathcal{E}_*] - f_{ij}^2 + (\text{Var}[d_{ij}/p] + f_{ij}^2) \Pr[\mathcal{E}_0] \end{aligned}$$

Note that

$$\begin{aligned} \text{Var}[\hat{f}_{ij}] &= \mathbb{E}[\hat{f}_{ij}^2] - f_{ij}^2 \\ &= \mathbb{E}[\hat{f}_{ij}^2 | \mathcal{E}_*] \Pr[\mathcal{E}_*] + \mathbb{E}[\hat{f}_{ij}^2 | \mathcal{E}_0] \Pr[\mathcal{E}_0] - f_{ij}^2 \\ &= \mathbb{E}[\hat{f}_{ij}^2 | \mathcal{E}_*] \Pr[\mathcal{E}_*] + f_{ij}^2 \Pr[\mathcal{E}_0] - f_{ij}^2, \end{aligned}$$

so

$$\begin{aligned} \text{Var}[\hat{f}'_{ij}] &= \text{Var}[\hat{f}_{ij}] + \text{Var}[d_{ij}/p] \Pr[\mathcal{E}_0] \\ &\leq \text{Var}[\hat{f}_{ij}] + \frac{f_{ij}}{p} \cdot (1-p)^{f_{ij}}. \end{aligned}$$

Due to the same reason as above, the second term is $O(1/p^2)$, and the proof completes. \square

Dealing with a decreasing p .

As in the count-tracking algorithm, we divide the whole tracking period into $O(\log N)$ rounds. Within each round, n stays within a constant factor of \bar{n} , while \bar{n} remains fixed for the whole round.

Within a round, we set the parameter p for all sites to be $p = 1/\lfloor \varepsilon \bar{n} / \sqrt{k} \rfloor_2$. When we proceed to a new round, all sites clear their memory and start a new copy of the algorithm from scratch with the new p . Given an item j , the coordinator estimates its frequency from each round separately, and add them up. Since the variance in a round is $O(k/p^2)$ and p increases geometrically over the rounds, the total variance is asymptotically bounded by the variance of the last round, i.e., $O(1/\varepsilon^2)$, as desired.

The space used at some site could still be large, since the site may receive too many elements in a round. If all the $O(n)$ elements in a round have gone to the same site, the site will need to use space $O(pn) = O(\sqrt{k}/\varepsilon)$. To bound the space, we restrict the amount of space used by each site. More precisely, when a site receives more than \bar{n}/k elements, it sends a message to the coordinator for notification, clears its memory, and starts a new copy of the algorithm from scratch. The coordinator will treat the new copy as if it were a new site, while the original site no longer receives more elements. Now the space used at each site is at most $p\bar{n}/k = O(1/(\varepsilon\sqrt{k}))$. Since there are at most $O(k)$ such new “virtual” sites ever created in a round, this does not affect the variance by more than a constant factor.

It remains to show that the total communication cost is $O(\sqrt{k}/\varepsilon \cdot \log N)$. From earlier we know that there are $O(\log N)$ rounds; within each round, \bar{n} is the same and n stays within $\Theta(\bar{n})$. Focus on one round. For each arriving element, the site S_i updates \bar{c}_{ij} with probability p and also independently samples it with probability p to maintain d_{ij} . This costs $O(n \cdot p) = O(\sqrt{k}/\varepsilon)$ communication.

THEOREM 3.1. *There is an algorithm for the frequency-tracking problem that, at any time, estimates the frequency of any element within error εn with probability at least 0.9. It uses $O(1/(\varepsilon\sqrt{k}))$ space at each site and $O(\sqrt{k}/\varepsilon \cdot \log N)$ communication.*

3.2 Space lower bound

It is easy to see that the communication lower bounds for the count-tracking problem also hold for the frequency-tracking problem. In this section, we prove the following space-communication trade-off.

THEOREM 3.2. *Consider any randomized algorithm for the frequency-tracking problem that, at any time, estimates the frequency of any element within error εn with probability at least 0.9. If the algorithm uses C bits of communication and uses M bits of space per site, then we must have $C \cdot M = \Omega(\log N/\varepsilon^2)$, assuming $k \leq 1/\varepsilon^2$.*

Thus, if the communication cost is $C = O(\sqrt{k}/\varepsilon \cdot \log N)$ bits, the space required per site is at least $\Omega(1/(\varepsilon\sqrt{k}))$ bits, as claimed in Table 1. If we ignore the word/bit difference, the space bounds are also tight. Interestingly, this lower

bound also shows that the random sampling algorithm [9] (see Table 1) actually attains the other end of this space-communication trade-off.

Proof. (of Theorem 3.2) We will use a result in [26] which states that, under the k -party communication model, there is an input distribution μ_k such that, any algorithm that solves the one-shot version of the problem under μ_k with error $2\varepsilon n$ with probability 0.9 needs at least $c\sqrt{k}/\varepsilon$ bits of communication for some constant c , assuming $k \leq 1/\varepsilon^2$. Moreover, any algorithm that solves ℓ independent copies of the one-shot version of the problem needs at least $\ell \cdot c\sqrt{k}/\varepsilon$ bits of communication.

We will consider the problem over ρk sites, for some integer $\rho \geq 1$ to be determined later. We divide the whole tracking period into $\log N$ rounds. In each round $i = 1, \dots, \log N$, we generate an input independently chosen from distribution $\mu_{\rho k}$ to the sites. We pick elements from a different domain for every round so that we have $\log N$ independent instances of the problem. In round i , for every element e picked from $\mu_{\rho k}$ for any site, we replace it with 2^{i-1} copies of e . We arrange the element arrivals in a round so that site S_1 gets all its elements first, then S_2 gets all its elements, and so on so forth. We will only require the continuous tracking algorithm to solve the frequency estimation problem at the end of each round. Since the last round always contains half of all the elements that have arrived so far, the algorithm must solve the problem for the elements in each round, namely, $\log N$ independent instances of the one-shot problem. By the result in [26], the communication cost to solve all these instances of the problem is at least $c\sqrt{\rho k}/\varepsilon \cdot \log N$.

Let \mathcal{A}_k be a continuous tracking algorithm over k sites that communicates C bits in total and uses M bits of space per site. Below we show how to solve the problem over the ρk sites in each round, by simulating the k -site algorithm \mathcal{A}_k . In each round, we start the simulation with sites S_1, \dots, S_k . Whenever \mathcal{A}_k exchanges a message, we do the same. When S_1 has received all its elements, it sends its memory content to S_{k+1} , which then takes the role of S_1 in the simulation and continues. Similarly, when S_2 has received all its elements, it sends its memory content to S_{k+2} , which replaces S_2 in the simulation. In general, when S_j is done with all its elements, it passes its role to S_{j+k} . When $S_{\rho k}$ is done, the simulation finishes for this round. $S_{\rho k}$ then sends a broadcast message and we proceed to the next round.

Let us analyze the communication cost of the simulation. First, we exchange exactly the same messages as \mathcal{A}_k does, which costs C . We also communicate $\rho(k-1)$ memory snapshots and a broadcast message in each round, which costs $\leq \rho k M \log N$ over all rounds. Thus, we have

$$C + \rho k M \log N \geq c\sqrt{\rho k}/\varepsilon \cdot \log N.$$

Rearranging,

$$M \geq \frac{c}{\varepsilon\sqrt{\rho k}} - \frac{C}{\rho k \log N} = \frac{1}{\sqrt{\rho k}} \left(\frac{c}{\varepsilon} - \frac{C}{\sqrt{\rho k} \log N} \right).$$

Thus, if we set $\sqrt{\rho} = \left\lceil \frac{2C\varepsilon}{c\sqrt{k} \log N} \right\rceil$, then

$$M \geq \frac{c}{2\varepsilon\sqrt{\rho k}} = \Omega\left(\frac{\log N}{C\varepsilon^2}\right),$$

as claimed. \square

4. TRACKING DISTRIBUTED RANKS

On a stream of n elements, an algorithm that produces an unbiased estimator for any rank with variance $O((\varepsilon n)^2)$ was presented in [22], which has been very recently improved and made to work in a stronger model [1]. It uses $O(1/\varepsilon \cdot \log^{1.5}(1/\varepsilon))$ working space to maintain a rank estimation summary structure of size $O(1/\varepsilon)$. We call this algorithm \mathcal{A} and will use it as a black box in our distributed tracking algorithm.

The overall algorithm.

As before, with $O(k \log N)$ communication, we first track \bar{n} , a constant factor approximation of the current n . This also divides the tracking period into $O(\log N)$ rounds. The $\Theta(n)$ elements arriving in a round are divided into chunks of size at most \bar{n}/k , each processed by an instance of algorithm \mathcal{C} , described below. A site may receive more than \bar{n}/k elements. When the $(\bar{n}/k + 1)$ th element arrives, the site finishes the current instance of \mathcal{C} , and starts a new one, which will process the next \bar{n}/k elements, and so on so forth.

Algorithm \mathcal{C} .

Algorithm \mathcal{C} reads at most \bar{n}/k elements, and divides them into blocks of size $b = \varepsilon \bar{n}/\sqrt{k}$, so there are at most $\frac{1}{\varepsilon \sqrt{k}}$ blocks. The algorithm builds a balanced binary tree on the blocks in the arrival order, and the height of the tree is $h \leq \log \frac{1}{\varepsilon \sqrt{k}}$. For each node v in the tree, let $D(v)$ be all the elements contained in the leaves in the subtree rooted at v . For each $D(v)$, the site starts an instance of \mathcal{A} , denoted as \mathcal{A}_v , to process its elements as they arrive. We say that v is *active* if \mathcal{A}_v is still accepting elements. For a node v at level ℓ (the leaves are said to be on level 0), the error parameter of \mathcal{A}_v is set to $2^{-\ell}/\sqrt{h}$. We say v is *full* if all the elements in $D(v)$ have arrived. When v is full, the site sends the summary computed by \mathcal{A}_v to the coordinator, and free the space used by \mathcal{A}_v . Furthermore, for each element that is arriving, the site samples it with probability $p = \frac{\sqrt{k}}{\varepsilon \bar{n}}$, and if it is sampled, the site sends it to the coordinator.

Analysis of costs.

We first analyze the various costs of \mathcal{C} . At any time there are at most h active nodes, one at each level, so the space used by \mathcal{C} is at most

$$\sum_{\ell=0}^h \sqrt{h} 2^\ell \log^{1.5} \frac{1}{\varepsilon} = O\left(\frac{\sqrt{h}}{\varepsilon \sqrt{k}} \log^{1.5} \frac{1}{\varepsilon}\right).$$

The communication for \mathcal{C} includes all the summaries computed, and the elements sampled. For each ℓ , the total size of the summaries on level ℓ is

$$O\left(\frac{1}{\varepsilon \sqrt{k}} 2^{-\ell} \cdot 2^\ell \sqrt{h}\right) = O\left(\frac{\sqrt{h}}{\varepsilon \sqrt{k}}\right).$$

Summing over all h levels, it is $\frac{h^{1.5}}{\varepsilon \sqrt{k}}$. There are at most $2k$ instances of \mathcal{C} in a round, therefore the total communication cost in a round is $O(h^{1.5} \sqrt{k}/\varepsilon)$. The number of sampled elements in a round is $O(np) = O(\sqrt{k}/\varepsilon)$. Thus, over all $O(\log N)$ rounds, the total communication cost is $O(h^{1.5} \sqrt{k}/\varepsilon \cdot \log N)$.

Estimation.

It remains to show how the coordinator estimates the rank of any given element x at any time with variance $O((\varepsilon n)^2)$. We decompose all n elements that have arrived so far into smaller subsets, and estimate the rank of x in each of the subsets. Since all estimators are unbiased, the overall estimator is also unbiased; the variance will be the sum of all the variances.

We will focus on the current round; all previous rounds can be handled similarly. Recall that there are $O(\bar{n})$ elements arriving in this round and $\bar{n} = \Theta(n)$. Every chunk of \bar{n}/k elements are processed by one instance of \mathcal{C} . Consider any such chunk. Suppose up to now, n' elements in this chunk have arrived for some $n' \leq \bar{n}/k$. We write n' as $n' = q \cdot b + r$ for some $r < b$, and decompose these n' elements into at most $h + 1$ subsets. The first qb elements are decomposed into at most h subsets, each of which corresponds to a full node in the binary tree of \mathcal{C} . The node has already sent its summary to the coordinator, which we can use to estimate the rank. For a node at level ℓ , the variance is $((2^{-\ell}/\sqrt{h}) \cdot 2^\ell b)^2 = b^2/h$, so the total variance from all h nodes is b^2 .

For the last r elements of the chunk that are still being processed by an active node, the coordinator does not have any summary for them. But recall that the site always samples each element with probability $p = \sqrt{k}/(\varepsilon \bar{n})$ and sends it to the coordinator if it is sampled. Thus, the rank of x in these r elements can be estimated by simply counting the number c of elements sampled that are smaller than x , and the estimator is c/p . The variance of this estimator is $r/p \leq b/p = b^2$. Thus, the variance from any chunk is $O(b^2)$. Since there are at most $2k$ chunks in the round, the total variance is $O(b^2 k) = O((\varepsilon \bar{n})^2) = O((\varepsilon n)^2)$. As the variances of the previous rounds are geometrically decreasing, the total variance from all the rounds is still bounded by $O((\varepsilon n)^2)$, as desired.

THEOREM 4.1. *There is an algorithm for the rank-tracking problem that, at any time, estimate the rank of any element within error εn with probability at least 0.9. It uses $O\left(\frac{1}{\varepsilon \sqrt{k}} \log^{1.5} \frac{1}{\varepsilon} \log^{0.5} \frac{1}{\varepsilon \sqrt{k}}\right)$ space at each site with communication cost $O\left(\frac{\sqrt{k}}{\varepsilon} \log N \log^{1.5} \frac{1}{\varepsilon \sqrt{k}}\right)$.*

5. REFERENCES

- [1] P. K. Agarwal, G. Cormode, Z. Huang, J. M. Phillips, Z. Wei, and K. Yi. Mergeable summaries. In *Proc. ACM Symposium on Principles of Database Systems*, 2012.
- [2] C. Arackaparambil, J. Brody, and A. Chakrabarti. Functional monitoring without monotonicity. In *Proc. International Colloquium on Automata, Languages, and Programming*, 2009.
- [3] B. Babcock and C. Olston. Distributed top-k monitoring. In *Proc. ACM SIGMOD International Conference on Management of Data*, 2003.
- [4] Z. Bar-Yossef. *The complexity of massive data set computations*. PhD thesis, University of California at Berkeley, 2002.
- [5] H.-L. Chan, T. W. Lam, L.-K. Lee, and H.-F. Ting. Continuous monitoring of distributed data streams over a time-based sliding window. *Algorithmica*, 62(3–4):1088–1111, 2011.




Figure 1: Differentiating two distributions

- [6] G. Cormode, M. Garofalakis, S. Muthukrishnan, and R. Rastogi. Holistic aggregates in a networked world: Distributed tracking of approximate quantiles. In *Proc. ACM SIGMOD International Conference on Management of Data*, 2005.
- [7] G. Cormode and M. Hadjieleftheriou. Finding frequent items in data streams. In *Proc. International Conference on Very Large Data Bases*, 2008.
- [8] G. Cormode, S. Muthukrishnan, and K. Yi. Algorithms for distributed functional monitoring. *ACM Transactions on Algorithms*, 7(2), Article 21, 2011. Preliminary version in SODA'08.
- [9] G. Cormode, S. Muthukrishnan, K. Yi, and Q. Zhang. Continuous sampling from distributed streams. *Journal of the ACM*, 59(2), 2012. Preliminary version in PODS'10.
- [10] W. Feller. *An introduction to probability theory and its applications*. Wiley, New York, 1968.
- [11] P. B. Gibbons and S. Tirthapura. Estimating simple functions on the union of data streams. In *Proc. ACM Symposium on Parallelism in Algorithms and Architectures*, 2001.
- [12] M. Greenwald and S. Khanna. Space-efficient online computation of quantile summaries. In *Proc. ACM SIGMOD International Conference on Management of Data*, 2001.
- [13] Z. Huang, L. Wang, K. Yi, and Y. Liu. Sampling based algorithms for quantile computation in sensor networks. In *Proc. ACM SIGMOD International Conference on Management of Data*, 2011.
- [14] Z. Huang, K. Yi, Y. Liu, and G. Chen. Optimal sampling algorithms for frequency estimation in distributed data. In *IEEE INFOCOM*, 2011.
- [15] R. Keralapura, G. Cormode, and J. Ramamirtham. Communication-efficient distributed monitoring of thresholded counts. In *Proc. ACM SIGMOD International Conference on Management of Data*, 2006.
- [16] A. Manjhi, V. Shkapenyuk, K. Dhamdhere, and C. Olston. Finding (recently) frequent items in distributed data streams. In *Proc. IEEE International Conference on Data Engineering*, 2005.
- [17] G. Manku and R. Motwani. Approximate frequency counts over data streams. In *Proc. International Conference on Very Large Data Bases*, 2002.
- [18] A. Metwally, D. Agrawal, and A. Abbadi. An integrated efficient solution for computing frequent and top-k elements in data streams. *ACM Transactions on Database Systems*, 31(3):1095–1133, 2006.
- [19] J. Misra and D. Gries. Finding repeated elements. *Science of Computer Programming*, 2:143–152, 1982.
- [20] J. I. Munro and M. S. Paterson. Selection and sorting with limited storage. *Theoretical Computer Science*, 12:315–323, 1980.
- [21] B. Patt-Shamir and A. Shafrir. Approximate distributed top-k queries. *Distributed Computing*, 21(1):1–22, 2008.
- [22] S. Suri, C. Toth, and Y. Zhou. Range counting over multidimensional data streams. *Discrete and Computational Geometry*, 2006.
- [23] S. Tirthapura and D. P. Woodruff. Optimal random sampling from distributed streams revisited. In *Proc. International Symposium on Distributed Computing*, 2011.
- [24] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971.
- [25] D. P. Woodruff. *Efficient and Private Distance Approximation in the Communication and Streaming Models*. PhD thesis, Massachusetts Institute of Technology, 2007.
- [26] D. P. Woodruff and Q. Zhang. Tight bounds for distributed functional monitoring. In *Proc. ACM Symposium on Theory of Computing*, 2012.
- [27] A. C. Yao. Probabilistic computations: Towards a unified measure of complexity. In *Proc. IEEE Symposium on Foundations of Computer Science*, 1977.
- [28] K. Yi and Q. Zhang. Optimal tracking of distributed heavy hitters and quantiles. In *Proc. ACM Symposium on Principles of Database Systems*, 2009.

APPENDIX

A. LOWER BOUND FOR THE SAMPLING PROBLEM

CLAIM A.1. *To solve the sampling problem we need to probe at least $\Omega(k)$ sites.*

Proof. Suppose that the coordinator only samples $z = o(k)$ sites. Let X be the number of sites that are sampled with bit 1. Then X is chosen from the hypergeometric distribution with probability density function (pdf) $\Pr[X =$

$x] = \binom{s'}{x} \binom{k'-s'}{z-x} / \binom{k'}{z}$. The expected value of X is $\frac{z}{k'} \cdot s'$, which is $\frac{z}{k'} \left(\frac{k}{2} - y + \sqrt{k} \right)$ or $\frac{z}{k'} \left(\frac{k}{2} - y - \sqrt{k} \right)$, depending on the value of s' . Let $p = \left(\frac{k}{2} - y \right) / k' = \frac{1}{2} \pm o(1)$ and $\alpha = \sqrt{k}/k' = 1/\sqrt{k} \pm o(1/\sqrt{k})$. To avoid tedious calculation, we assume that X is picked randomly from one of the two normal distributions $\mathcal{N}_1(\mu_1, \sigma_1^2)$ and $\mathcal{N}_2(\mu_2, \sigma_2^2)$ with equal probability, where $\mu_1 = z(p - \alpha)$, $\mu_2 = z(p + \alpha)$, $\sigma_1, \sigma_2 = \Theta(\sqrt{zp(1-p)}) = \Theta(\sqrt{z})$. In Feller [10] it is shown that the normal distribution approximates the hypergeometric distribution very well when z is large and $p \pm \alpha$ are constants in $(0, 1)$ ⁶. Now our task is to decide from which of the two distributions X is drawn based on the value of X with success probability at least 0.7.

Let $f_1(x; \mu_1, \sigma_1^2)$ and $f_2(x; \mu_2, \sigma_2^2)$ be the pdf of the two normal distributions $\mathcal{N}_1, \mathcal{N}_2$, respectively. It is easy to see that the best deterministic algorithm of differentiating the two distributions based on the value of a sample X will do the following.

- If $X > x_0$, then X is chosen from \mathcal{N}_2 , otherwise X is chosen from \mathcal{N}_1 , where x_0 is the value such that $f_1(x_0; \mu_1, \sigma_1^2) = f_2(x_0; \mu_2, \sigma_2^2)$ (thus $\mu_1 < x_0 < \mu_2$).

Indeed, if $X > x_0$ and the algorithm decides that “ X is chosen from \mathcal{N}_1 ”, we can always flip this decision and improve the success probability of the algorithm.

The error comes from two sources: (1) $X > x_0$ but X is actually drawn from \mathcal{N}_2 ; (2) $X \leq x_0$ but X is actually drawn from \mathcal{N}_1 . The total error is

$$1/2 \cdot (\Phi(-\ell_1/\sigma_1) + \Phi(-\ell_2/\sigma_2)),$$

where $\ell_1 = x_0 - \mu_1$ and $\ell_2 = \mu_2 - x_0$. (Thus $\ell_1 + \ell_2 = \mu_2 - \mu_1 = 2\alpha z$). $\Phi(\cdot)$ is the cumulative distribution function (cdf) of the normal distribution. See Figure 1.

Finally note that $\ell_1/\sigma_1 = O(\alpha z/\sqrt{z}) = O(\sqrt{z/k}) = o(1)$ and $\ell_2/\sigma_2 = O(\alpha z/\sqrt{z}) = o(1)$, so $\Phi(-\ell_1/\sigma_1) + \Phi(-\ell_2/\sigma_2) > 0.99$. Therefore, the failure probability is at least 0.49, contradicting our success probability guarantee. Thus we must have $z = \Omega(k)$. \square

⁶In Feller’s book [10] the following is proved. Let $p \in (0, 1)$ be some constant and $q = 1 - p$. The population size is N and the sample size is n , so that $n < N$ and Np, Nq are both integers. The hypergeometric distribution is $P(k; n, N) = \binom{Np}{k} \binom{Nq}{n-k} / \binom{N}{n}$ for $0 \leq k \leq n$.

THEOREM A.1. [10] *If $N \rightarrow \infty, n \rightarrow \infty$ so that $n/N \rightarrow t \in (0, 1)$ and $x_k := (k - np)/\sqrt{npq} \rightarrow x$, then*

$$p(k; n, N) \sim \frac{e^{-x^2/2(1-t)}}{\sqrt{2\pi npq(1-t)}}$$