On Competitive Algorithms for Approximations of Top-*k*-Position Monitoring of Distributed Streams

Alexander Mäcker Manuel Malatyali Friedhelm Meyer auf der Heide

Heinz Nixdorf Institute & Computer Science Department Paderborn University, Germany {amaecker, malatya, fmadh}@hni.upb.de

Abstract

Consider the continuous distributed monitoring model in which n distributed nodes, receiving individual data streams, are connected to a designated server. The server is asked to continuously monitor a function defined over the values observed across all streams while minimizing the communication. We study a variant in which the server is equipped with a broadcast channel and is supposed to keep track of an approximation of the set of nodes currently observing the k largest values. Such an approximate set is exact except for some imprecision in an ε -neighborhood of the k-th largest value. This approximation of the Top-k-Position Monitoring Problem is of interest in cases where marginal changes (e.g. due to noise) in observed values can be ignored so that monitoring an approximation is sufficient and can reduce communication.

This paper extends our results from [6], where we have developed a filter-based online algorithm for the (exact) Top-k-Position Monitoring Problem. There we have presented a competitive analysis of our algorithm against an offline adversary that also is restricted to filter-based algorithms. Our new algorithms as well as their analyses use new methods. We analyze their competitiveness against adversaries that use both exact and approximate filter-based algorithms, and observe severe differences between the respective powers of these adversaries.

1 Introduction

We consider a setting in which n distributed nodes are connected to a central server. Each node continuously observes a data stream and the server is asked to keep track of the value of some function defined over all streams. In order to fulfill this task, nodes can communicate to the server, while the server can employ a broadcast channel to send a message to all nodes.

In an earlier paper [6], we introduced and studied a problem called Top-k-Position Monitoring in which, at any time t, the server is interested in monitoring the k nodes that are observing the largest values at this particular time t. As a motivating example, picture a scenario in which a central load balancer within a local cluster of webservers is interested in keeping track of those nodes which are facing the highest loads. We proposed an algorithm based on the notion of filters and analyzed its competitiveness with respect to an optimal filter-based offline algorithm. Filters are assigned by the server and are used as a means to indicate the nodes when they can resign to send updates; this particularly reduces communication when observed values are "similar" to the values observed in the previous time steps.

In this paper, we broaden the problem and investigate the monitoring of an approximation of the Top-*k*-Positions. We study the problem of ε -Top-*k*-Position Monitoring, in which the server is supposed to maintain a subset of *k* nodes such that all nodes observing "clearly larger" values than the node which observed the *k*-th largest value are within this set and no node observing a "clearly smaller" value belongs to this set. Here, smaller/larger is meant to be understood with respect to ε and the *k*-th largest value observed. A detailed definition is given in Sect. 2. Relaxing the problem in this direction can reduce communication while, in many cases, marginal or insignificant changes (e.g. due to noise)

This work was partially supported by the German Research Foundation (DFG) within the Priority Program "Algorithms for Big Data" (SPP 1736) and by the EU within FET project MULTIPLEX under contract no. 317532.

in observed values can be ignored and justify the sufficiency of an approximation. Examples are situations where lots of nodes observe values oscillating around the *k*-th largest value and where this observation is not of any qualitative relevance for the server. We design and analyze algorithms for ε -Top-*k*-Position Monitoring and, although we use these very tools of filters and competitive analysis [6], the imprecision/approximation requires fundamentally different online strategies for defining filters in order to obtain efficient solutions.

1.1 Our Contribution

In this paper we investigate a class of algorithms that are based on using filters and study their efficiency in terms of competitive analysis.

As a first technical contribution we analyze an algorithm (Sect. 3) which allows the server to decide the logical disjunction of the (binary) values observed by the distributed nodes. It uses a logarithmic number of rounds and a constant number of messages on expectation. As a by-product, using this algorithm, the result on the competitiveness of the filter-based online algorithm in [6] can be reduced from $O(k \log n + \log \Delta \log n)$ to $O(k \log n + \log \Delta)$, for observed values from $\{0, 1, \ldots, \Delta\}$.

Second, we also propose an online algorithm (Sect. 4) that is allowed to introduce an error of $\varepsilon \in (0, 1/2]$ in the output and compare it to an offline algorithm that solves the exact Top-*k*-Position Monitoring problem. We show that this algorithm is $\mathcal{O}(k \log n + \log \log \Delta + \log \frac{1}{\varepsilon})$ -competitive. Note that this imprecision allows to bring the $\log \Delta$ in the upper bound down to $\log \log \Delta$ for any constant ε .

We also investigate the setting in which also the offline algorithm is allowed to have an error in the output (Sect. 5). We first show that these results are not comparable to previous results; we prove a lower bound on the competitiveness of $\Omega(n/k)$. Our third and main technical contribution is an algorithm with a competitiveness of $\mathcal{O}(n^2 \log(\varepsilon \Delta) + n \log^2(\varepsilon \Delta) + \log \log \Delta + \log \frac{1}{\varepsilon})$ if the online and the offline algorithm may use an error of ε .

However, if we slightly decrease the allowed error for the offline algorithm, the lower bound on the competitiveness of $\Omega(n/k)$ still holds, while the upper bound is reduced to $\mathcal{O}(n+k\log n + \log\log \Delta + \log \frac{1}{\epsilon})$.

1.2 Related Work

Efficient computation of functions on big datasets in terms of streams has turned out to be an important topic of research with applications in network traffic analysis, text mining or databases (e.g. [9] and [7]).

The *Continuous Monitoring Model*, which we consider in this paper, was introduced by Cormode et al. [2] to model systems comprised of a server and n nodes observing *distributed* data streams. The primary goal addressed within this model is the continuous computation of a function depending on the information available across all n data streams up to the current time at a dedicated server. Subject to this main concern, the minimization of the overall number of messages exchanged between the nodes and the server usually determines the efficiency of a streaming algorithm. We refer to this model and enhance it by a broadcast channel as proposed by Cormode et al. in [3].

An important class of problems investigated in literature are threshold computations where the server is supposed to decide whether the current function value has reached some given threshold τ . For monotone functions such as monitoring the number of distinct values or the sum over all values, exact characterizations in the deterministic case are known [2, 3]. However, non-monotone functions, e.g., the entropy [1], turned out to be much more complex to handle.

A general approach to reduce the communication when monitoring distributed streams is proposed in [12]. Zhang et al. introduce the notion of *filters*, which are also an integral part of our algorithms. They consider the problem of continuous skyline maintenance, in which a server is supposed to continuously maintain the skyline of dynamic objects. As they aim at minimizing the communication overhead between the server and the objects, they use a filter method that helps in avoiding the transmission of updates in case these updates cannot influence the skyline. More precisely, the objects are points of a *d*-dimensional space and filters are hyper-rectangles assigned by the server to the objects such that as long as these points are within the assigned hyper-rectangle, updates need not be communicated to the server.

Despite its online nature, by now streaming algorithms are barely studied in terms of competitiveness. In their work [11], Yi and Zhang were the first to study streaming algorithms with respect to their competitiveness and recently this approach was also applied in a few papers ([5, 10, 6, 4]). In their model [11], there is one node and one server and the goal is to keep the server informed about the current value of a function $f : \mathbb{Z}^+ \to \mathbb{Z}^d$ that is observed by the node and changes its value over time, while minimizing the number of messages. Yi and Zhang present an algorithm that

is $\mathcal{O}(d^2 \log(d \cdot \delta))$ -competitive if the last value received by the server might deviate by δ from the current value of f. Recently, Tang et al. [10] extended this work by Yi and Zhang for the two-party setting to the distributed case. They consider a model in which the server is supposed to track the current value of a (one-dimensional) function that is defined over a set of n functions observed at the distributed nodes. Among other things, they propose an algorithm for the case of a tree-topology in which the distributed nodes are the leaves of a tree connecting them to the server. They show that on any instance I their algorithm incurs communication cost that is by a factor of $\mathcal{O}(h_{max} \log \delta)$, where h_{max} represents the maximinum length of a path in the tree, larger than those of the best solution obtained by an online algorithm on I.

Following the idea of studying competitive algorithms for monitoring streams and the notion of filters, Lam et al. [5] present an algorithm for online dominance tracking of distributed streams. In this problem a server always has to be informed about the dominance relationship between n distributed nodes each observing an online stream of d-dimensional values. Their algorithm is based on the idea of filters and they show that a mid-point strategy, which sets filters to be the mid-point between neighboring nodes, is $O(d \log U)$ -competitive with respect to the number of messages sent in comparison to an offline algorithm that sets filters optimally.

While we loosely motivated our search for approximate solutions by noise in the introduction, in other problems noise is a major concern and explicitly addressed. For example, consider streaming algorithms for estimating statistical parameters like frequency moments [13]. In such problems, certain elements from the universe may appear in different forms due to noise and thus, should actually be treated as the same element.

2 Preliminaries

In our setting there are *n* distributed nodes $\{1, \ldots, n\}$. Each node *i* receives a continuous data stream $(v_i^1, v_i^2, v_i^3 \ldots)$, which can be exclusively observed by node *i*. At time *t*, $v_i^t \in \mathbb{N}$ is observed and no $v_i^{t'}$, t' > t, is known. We omit the index *t* if it is clear from the context.

Following the model in [3], we allow that between any two consecutive time steps, a *communication protocol* exchanging messages between the server and the nodes may take place. The communication protocol is allowed to use an amount of rounds which is polylogarithmic in n and $\max_{1 \le i \le n}(v_i^t)$. The nodes can communicate to the server while the server can communicate to single nodes or utilize a broadcast channel to communicate a message that is received by all nodes at the same time. These communication methods incur unit communication cost per message, we assume instant delivery, and a message at time t is allowed to have a size at most logarithmic in n and $\max_{1 \le i \le n}(v_i^t)$.

Problem Description Consider the Top-k-Position Monitoring problem [6], in which the server is asked to keep track of the set of nodes currently holding the k largest values. We relax this definition and study an approximate variant of the problem in which this set is exact except for nodes in a small neighborhood around the k-th largest value. We denote by $\pi(k,t)$ the node which observes the k-th largest value at time t and denote by top- $k := \{i \in \{1, \ldots, k\} : \pi(i, t)\}$ the nodes observing the k largest values. Given an error $0 < \varepsilon < 1$, for a time t we denote by $E(t) := (\frac{1}{1-\varepsilon}v_{\pi(k,t)}^{t}, \infty)$ the range of values that are clearly larger than the k-th largest value and by $A(t) := [(1-\varepsilon)v_{\pi(k,t)}^{t}, \frac{1}{1-\varepsilon}v_{\pi(k,t)}^{t}]$ the ε -neighborhood around the k-th largest value. Furthermore, we denote by $\mathcal{K}(t) := \{i : v_i^t \in A(t)\}$ the nodes in the ε -neighborhood around the k-th largest value. Then, at any time t, the server is supposed to know the nodes $\mathcal{F}(t) = \mathcal{F}_E(t) \cup \mathcal{F}_A(t) = \{i_1, \ldots, i_k\}$ according to the following properties:

1.
$$\mathcal{F}_{E}(t) = \{i : v_{i}^{t} \in E(t)\}$$
 and

2.
$$\mathcal{F}_A(t) \subseteq \mathcal{K}(t) = \{i : v_i^t \in A(t)\}$$
, such that $|\mathcal{F}_A(t)| = k - |\mathcal{F}_E(t)|$ holds

Denote by Δ the maximal value observed by some node (which may not be known beforehand). We use $\mathcal{F}_1 = \mathcal{F}(t)$ if t is clear from the context, $\mathcal{F}_2 = \{1, \ldots, n\} \setminus \mathcal{F}(t)$, and call \mathcal{F}^* the *output* of an optimal offline algorithm. If the k-th and the (k + 1)-st largest value differ by more than $\varepsilon v_{\pi(k,t)}^t$, $\mathcal{F}(t)$ coincides with the set in the (exact) Top-k-Position Monitoring problem and hence, $\mathcal{F}(t)$ is unique. We denote by $\sigma(t) := |\mathcal{K}(t)|$ the number of nodes at time t which are in the ε -neighborhood of the k-th largest value and $\sigma := \max_t \sigma(t)$. Note that $|\mathcal{K}(t)| = 1$ implies that $\mathcal{F}(t)$ is unique. Furthermore for solving the exact Top-k-Position Monitoring problem we assume that the values are distinct (at least by using the nodes' identifiers to break ties in case the same value is observed by several nodes).

2.1 Filter-Based Algorithms & Competitive Analysis

A set of filters is a collection of intervals, one assigned to each node, such that as long as the observed values at each node are within its respective interval, the output $\mathcal{F}(t)$ need not change. For the problem at hand, this general idea of filters translates to the following definition.

Definition 2.1. [6] For a fixed time t, a set of filters is defined as an n-tuple of intervals (F_1^t, \ldots, F_n^t) , $F_i \subseteq \mathbb{N} \cup \{\infty\}$ and $v_i \in F_i$, such that as long as the value of node i only changes within its interval (i.e. $v_i \in F_i$), the value of the output \mathcal{F} need not change.

Observe that each pair of filters (F_i, F_j) of nodes $i \in \mathcal{F}(t)$ and $j \notin \mathcal{F}(t)$ must be disjoint except for a small overlapping. This observation can be stated formally as follows.

Observation 2.2. For a fixed time t, an n-tuple of intervals is a set of filters if and only if for all pairs $i \in \mathcal{F}(t)$ and $j \notin \mathcal{F}(t)$ the following holds: $v_i \in F_i = [\ell_i, u_i], v_j \in F_j = [\ell_j, u_j]$ and $\ell_i \ge (1 - \varepsilon)u_j$.

In our model, we assume that nodes are assigned such filters by the server. If a node observes a value that is larger than the upper bound of its filter, we say the node *violates its filter from below*. A violation *from above* is defined analogously. If such a violation occurs, the node may report it and its current value to the server. In contrast to [6], we allow the server to assign "invalid" filters, i.e., there are affected nodes that directly observe a filter-violation. However, for such an algorithm to be correct, we demand that the intervals assigned to the nodes at the end of the protocol at time t and thus, before observations at time t + 1, constitute a (valid) set of filters. We call such an algorithm *filter-based*. Note that the fact that we allow invalid filters (in contrast to [6]) simplifies the presentation of the algorithms in the following. However, using a constant overhead the protocols can be changed such that only (valid) filters are sent to the nodes.

Competitiveness To analyze the quality of our online algorithms, we use analysis based on competitiveness and compare the communication induced by the algorithms to that of an adversary's offline algorithm.

Similar to [5] and [6], we consider adversaries that are restricted to use filter-based offline algorithms and hence, OPT is lower bounded by the number of filter updates. However, we compare our algorithms against several adversaries which differ in terms of whether their offline algorithm solves the exact Top-*k*-Position Monitoring Problem or ε -Top-*k*-Position Monitoring. The adversaries are assumed to be adaptive, i.e., values observed by a node are given by an adversary who knows the algorithm's code, the current state of each node and the server and the results of random experiments.

An online algorithm is said to have a competitiveness of c if the number of messages is at most by a factor of c larger than that of the adversary's offline algorithm.

2.2 Observations and Lemmas

Define for some fixed set $S \subseteq \{1, ..., n\}$ the minimum of the values observed by nodes in S during a time period [t, t'] as $MIN_S(t, t')$ and the maximum of the values observed during the same period as $MAX_S(t, t')$.

Definition 2.3. Let t, t' be given times with $t' \ge t$. For a subset of nodes $S \subseteq \{1, \ldots, n\}$ the values $MAX_S(t, t') \coloneqq \max_{t \le t^* \le t'} \max_{i \in S} (v_i^{t^*})$ and $MIN_S(t, t')$ are defined analogously.

Observe that it is sufficient for an optimal offline algorithm to only make use of two different filters F_1 and F_2 .

Proposition 2.4. Without loss of generality, we may assume that an optimal offline algorithm only uses two different filters at any time.

Proof. Let [t, t'] be an interval during which OPT does not communicate. We fix its output \mathcal{F}_1^* and define $\mathcal{F}_2^* := \{1, \ldots, n\} \setminus \mathcal{F}_1^*$. If OPT only uses two different filters throughout the interval, we are done. Otherwise, using \mathcal{F}_1^* as output throughout the interval [t, t'] and filters $F_1 = [MIN_{\mathcal{F}_1^*}(t, t'), \infty]$ and $F_2 = [0, MAX_{\mathcal{F}_2^*}(t, t')]$, which must be feasible due to the assumption that OPT originally assigned filters that lead to no communication, leads to no communication within the considered interval.

The following lemma generalizes a lemma in [6] to ε -Top-k-Position Monitoring. Assuming the optimal offline algorithm did not change the set of filters during a time period [t, t'], the minimum value observed by nodes in \mathcal{F}_1^* can only be slightly smaller than the maximum value observed by nodes in \mathcal{F}_2^* .

Lemma 2.5. If OPT uses the same set of filters F_1 , F_2 during [t, t'], then it holds $\operatorname{MIN}_{\mathcal{F}_1^*}(t, t') \ge (1-\varepsilon) \operatorname{MAX}_{\mathcal{F}_2^*}(t, t')$.

Proof. Assume to the contrary that OPT uses the same set of filters throughout the interval [t, t'] and outputs \mathcal{F}_1^* , but $\operatorname{MIN}_{\mathcal{F}_1^*}(t, t') < (1 - \varepsilon)\operatorname{MAX}_{\mathcal{F}_2^*}(t, t')$ holds. Then there are two nodes, $i \in \mathcal{F}_1^*$ and $j \notin \mathcal{F}_1^*$, and two times $t_1, t_2 \in [t, t']$, such that $v_i^{t_1} = \operatorname{MIN}_{\mathcal{F}_1^*}(t, t')$ and $v_j^{t_2} = \operatorname{MAX}_{\mathcal{F}_2^*}(t, t')$. Due to the definition of a set of filters and the fact that OPT has not communicated during [t, t'], OPT must have set the filter for node i to $[s_1, \infty]$, $s_1 \leq v_i^{t_1}$, and for node j to $[-\infty, s_2]$, $s_2 \geq v_i^{t_2}$. This is a contradiction to the definition of a set of filters and Observation 2.2.

At last a result from [6] is restated in order to calculate the (exact) top-k set for one time step.

Lemma 2.6. [6] There is an algorithm that computes the node holding the largest value using $O(\log n)$ messages on expectation.

3 Auxiliary Problem: Existence

In our competitive algorithms designed and analyzed in the following, we will frequently make use of a protocol for a subproblem which we call EXISTENCE: Assume all nodes observe only binary values, i.e. $\forall i \in \{1, ..., n\} : v_i \in \{0, 1\}$. The server is asked to decide the *logical disjunction* for one fixed time step t.

It is known that for n nodes each holding a bit vector of length m the communication complexity to decide the bit-wise disjunction is $\Omega(nm)$ in the server model [8]. Observe that in our model 1 message is sufficient to decide the problem assuming the nodes have a unique identifier between 1 and n and the protocol uses n rounds.

We prove that it is sufficient to use a constant amount of messages on expectation and logarithmic number of rounds. Note that the algorithm in the following lemma is a Las Vegas algorithm, i.e. the algorithm is always correct and the number of messages needed is based on a random process.

Lemma 3.1. There is an algorithm EXISTENCEPROTOCOL that uses O(1) messages on expectation to solve the problem EXISTENCE.

Proof. Initially all nodes are active. All nodes *i* deactivate themselves, if $v_i = 0$ holds, that is, these nodes do not take part in the following process. In each round $r = 0, 1, ..., \log n$ the active nodes send messages independently at random with probability $p_r := 2^r/n$. Consequently, if the last round $\gamma = \log n$ is reached, all active nodes *i* with $v_i = 1$ send a message with probability 1. As soon as at least one message was sent or the γ -th round ends, the protocol is terminated and the server can decide EXISTENCE.

Next, we analyze the above protocol and show that the bound on the expected number of messages is fulfilled. Let X be the random variable for the number of messages used by the protocol and b be the number of nodes i with $v_i = 1$. Note that the expected number of messages sent in round r is $b \cdot p_r$ and the probability that no node has sent a message before is $\prod_{k=0}^{r-1} (1-p_k)^b$.

Observing that the function $f(r) = b \cdot p_r (1 - p_{r-1})^b$ has only one extreme point and $0 \le f(r) < 2$ for $r \in [0, \log n]$, it is easy to verify that the series can be upper bounded by simple integration:

$$\begin{split} \mathbb{E}[X] &\leq \frac{b}{n} + \sum_{r=1}^{\log(n)} \frac{b2^r}{n} \prod_{k=0}^{r-1} \left(1 - \frac{2^k}{n}\right)^b \\ &\leq 1 + \sum_{r=1}^{\log(n)} \frac{b2^r}{n} \left(1 - \frac{2^{r-1}}{n}\right)^b \\ &\leq 1 + \int_0^{\log(n)} \frac{b2^r}{n} \left(1 - \frac{2^{r-1}}{n}\right)^b dr + 2 \end{split}$$

$$\begin{split} &\leq 3 + \left[\frac{b}{(b+1)n\ln(2)}(2^r - 2n)\left(1 - \frac{2^{r-1}}{n}\right)^b\right]_0^{\log n} \\ &\leq 3 + \frac{1}{n\ln(2)} \cdot \\ &\left((2^{\log n} - 2n)\left(1 - \frac{2^{\log n-1}}{n}\right)^b + 2n\left(1 - \frac{2^{0-1}}{n}\right)^b\right) \\ &\leq 3 + \frac{1}{n\ln(2)}\left[(n-2n)\left(1 - \frac{1}{2}\right)^b + 2n\left(1 - \frac{1}{2n}\right)^b\right] \\ &\leq 3 + \frac{1}{n\ln(2)}\left[(-n)\frac{1}{2^b} + 2n\left(1 - \frac{1}{2n}\right)^b\right] \\ &\leq 3 + \frac{1}{\ln(2)}\left(2\left(1 - \frac{1}{2n}\right)^b - \frac{1}{2^b}\right) \\ &\leq 3 + \frac{1}{\ln(2)}\left(2 - \frac{1}{2^b}\right) \leq 3 + \frac{2}{\ln(2)} \leq 6 \end{split}$$

This protocol can be used for a variety of subtasks, e.g. validating that all nodes are within their filters, identifying that there is some filter-violation or whether there are nodes that have a higher value than a certain threshold.

Corollary 3.2. Given a time t. There is an algorithm which decides whether there are nodes which observed a filter-violation using O(1) messages on expectation.

Proof. For the distributed nodes to report filter-violations we use an approach based on the EXISTENCEPROTOCOL to reduce the number of messages sent in case several nodes observe filter-violations at the same time. The nodes apply the EXISTENCEPROTOCOL as follows: Each node that is still within its filter applies the protocol using a 0 as its value and each node that observes a filter-violation uses a 1. Note that by this approach the server definitely gets informed if there is some filter-violation and otherwise no communication takes place.

The EXISTENCEPROTOCOL can be used in combination with the relaxed definition of filters to strengthen the result for Top-k-Position Monitoring from $\mathcal{O}(k \log n + \log \Delta \log n)$ to $\mathcal{O}(k \log n + \log \Delta)$. We first introduce a generic framework and then show how to achieve this bound.

A generic approach Throughout the paper, several of our algorithms feature similar structural properties in the sense that they can be defined within a common framework. Hence, we now define a generic approach to describe the calculation and communication of filters, which we then refine later. The general idea is to only use two different filters that are basically defined by one value separating nodes in $\mathcal{F}(t)$ from the remaining nodes. Whenever a filter-violation is reported, this value is recalculated and used to set filters properly.

The approach proceeds in rounds. In the first round we define an initial interval L_0 . In the r-th round, based on interval L_r , we compute a value m that is broadcasted and is used to set the filters to [0, m] and $[m, \infty]$. As soon as node i reports a filter-violation observing the value v_i , the coordinator redefines the interval $L_{r+1} := L_r \cap [0, v_i]$ if the violation is from above and $L_{r+1} := L_r \cap [v_i, \infty]$ otherwise. The approach finishes as soon as some (predefined) condition is satisfied.

Corollary 3.3. There is an algorithm that is $O(k \log n + \log \Delta)$ -competitive for (exact) Top-k-Position Monitoring.

Proof. Our algorithm proceeds in phases that are designed such that we can show that an optimal algorithm needs to communicate at least once during a phase and additionally, we can upper bound the number of messages sent by the online algorithm according to the bound on the competitiveness.

We apply the generic approach with parameters described as follows. The initial interval is defined as $L_0 := [\ell, u]$, where $\ell = v_{\pi(k+1,t)}^t$, $u = v_{\pi(k,t)}^t$. This can be done by determining the values of the nodes holding the k + 1 largest values using $\mathcal{O}(k \log n)$ messages on expectation. In the *r*-th round, based on interval L_r , we compute the midpoint of

 L_r as the value *m* which is broadcasted and used to set the filters. As soon as a filter-violation is reported, the generic framework is applied. In case L_r is empty the phase ends.

Note that the distance between u and ℓ gets halved every time a node violates its filter leading to $\mathcal{O}(\log(u_0 - \ell_0)) = \mathcal{O}(\log \Delta)$ messages on expectation per phase. Also, it is not hard to see that during a phase *OPT* has communicated at least once and hence, we obtain the claimed bound on the competitiveness.

4 Competing against an Exact Adversary

In this section, we propose an algorithm based on the strategy to choose the nodes holding the k largest values as an output and use this set as long as it is feasible. It will turn out that this algorithm is suitable in two scenarios: First, it performs well against an adversary who solves the Top-k-Position Monitoring problem (cf. Theorem 4.5); second, we can use it in situations in which an adversary who is allowed to introduce some error and cannot exploit this error because the observed data leads to a unique output (cf. Sect. 5).

In particular, we develop an algorithm started at t that computes the output set $\mathcal{F}_1 := \mathcal{F}(t)$ using the protocol from Lemma 2.6 and for all consecutive times witnesses whether \mathcal{F}_1 is correct or not. Recall that while computing the set $\mathcal{F}(t)$ from scratch (cf. Lemma 2.6) is expensive in terms of communication, witnessing its correctness in consecutive rounds is cheap since it suffices to observe filter-violations (cf. Definition 2.1 and Corollary 3.2).

The algorithm tries to find a value m which partitions \mathcal{F}_1 from \mathcal{F}_2 according to the generic framework, such that for all nodes $i \in \mathcal{F}_1$ it holds $v_i \ge m$ and for all nodes $i \in \mathcal{F}_2$ it holds $v_i \le m$. We call such a value m certificate.

Guessing OPT's Filters In the following we consider a time period [t, t''] during which the output $\mathcal{F}(t)$ need not change. Consider a time $t' \in [t, t'']$. The online strategy to choose a certificate at this time contingents on the size of some interval

 L^* from which an offline algorithm must have chosen the lower bound ℓ^* of the upper filter at time t

such that the filters are valid throughout [t, t']. The algorithm TOP-K-PROTOCOL keeps track of (an approximation of) L^* at time t' denoted by $L = [\ell, u]$ for which $L^* \subseteq L$ holds. The online algorithm tries to improve the guess where OPT must have set filters by gradually reducing the size of interval L (while maintaining the invariant $L^* \subseteq L$) at times it observes filter-violations.

Initially u and ℓ are defined as follows: $u \coloneqq v_{\pi(k,t)}^t = \text{MIN}_{\mathcal{F}_1}(t,t)$ and $\ell \coloneqq v_{\pi(k+1,t)}^t = \text{MAX}_{\mathcal{F}_2}(t,t)$ and are redefined over time. Although defining the certificate as the midpoint of $L = [\ell, u]$ intuitively seems to be the best way to choose m, the algorithm is based on four consecutive phases, each defining a different strategy.

In detail, the first phase is executed as long as the property

$$\log\log u > \log\log \ell + 1 \tag{P1}$$

holds. In this phase, m is defined as $\ell + 2^{2^r}$ after r filter-violations observed. If the property

$$\log\log u \le \log\log \ell + 1 \land u > 4\ell \tag{P2}$$

holds, the value m is chosen to be 2^{mid} where mid is the midpoint of $[\log \ell, \log u]$. Observe that $2^{mid} \in L = [\ell, u]$ holds.

The third phase is executed if property

$$u \le 4 \ \ell \land u > \frac{1}{1 - \varepsilon} \ \ell \tag{P3}$$

holds and employs the intuitive approach of choosing m as the midpoint of L. The last phase contains the remaining case of

$$u \le \frac{1}{1 - \varepsilon} \ell \tag{P4}$$

and is simply executed until the next filter-violation is observed using the filters $F_1 = [\ell, \infty]$ and $F_2 = [0, u]$.

In the following we propose three algorithms A_1 , A_2 , and A_3 which are executed if the respective property hold and analyze the correctness and the amount of messages needed.

Lemma 4.1. Given time t, an output $\mathcal{F}(t)$, and an interval $L = [\ell, u]$ for which (P1) holds, there is an algorithm \mathcal{A}_1 that witnesses the correctness of $\mathcal{F}(t)$ until a time t' at which it outputs $L' = [\ell', u']$ for which (P1) does not hold. The algorithm uses $\mathcal{O}(\log \log \Delta)$ messages on expectation.

Proof. The algorithm \mathcal{A}_1 applies the generic framework and defines the value m, the server broadcasts, as $m := \ell_0 + 2^{2^r}$, where ℓ_0 is the initial value of ℓ . If $\log \log u' - \log \log \ell' \le 1$ holds, the algorithm terminates and outputs $L' = [\ell', u']$ with ℓ' and u' defined as the redefinition of ℓ and u respectively.

To analyze the amount of messages needed and express it in terms of Δ , observe that in the worst case the server only observes filter-violations from nodes $i \in \mathcal{F}_2$. In case there is a filter-violation from above, i.e. a node $i \in \mathcal{F}_1$ reports a filter-violation, the condition $\log \log u' - \log \log \ell' \le 1$ holds. At least in round $r = \log \log(u - \ell)$, which is by definition upper bounded by $\log \log \Delta$, the algorithm terminates.

If $\mathcal{F}(t)$ is not valid at time t', there are nodes $i_1 \in \mathcal{F}_1$, $i_2 \in \mathcal{F}_2$ and time points t_1, t_2 ($t_1 = t' \lor t_2 = t'$) for which $v_{i_1}^{t_1} < v_{i_2}^{t_2}$ holds. Thus, \mathcal{A}_1 observed a filter-violation by either i_1 or i_2 followed by a sequence alternating between filter-violations and filter-updates. At some point (but still at time t') $\log \log u' - \log \log \ell' \leq 1$ holds and the algorithm outputs (ℓ', u'), proving \mathcal{A}_1 's correctness for time t'.

Lemma 4.2. For a given $\mathcal{F}(t)$ and a given interval $L = [\ell, u]$ for which (P2) holds, there is an algorithm \mathcal{A}_2 that witnesses the correctness of $\mathcal{F}(t)$ until a time t' at which it outputs $L' = [\ell', u']$ for which (P2) does not hold. The algorithm uses $\mathcal{O}(1)$ messages on expectation.

Proof. We apply the generic approach and choose the value m to be broadcasted by 2^{mid} , where mid is the midpoint of $[\log \ell, \log u]$.

To analyze the amount of messages needed, bound $L = [\ell, u]$ in terms of values that are double exponential in 2. To this end, let $a \in \mathbb{N}$ be the largest number such that $\ell \geq 2^{2^a}$ holds. Now observe since (P2) holds, $u \leq 2^{2^{a+2}}$ follows. Since the algorithm chooses the midpoint of the interval $[\log \ell, \log u]$ in order to get m and halves this interval after every filter-violation, one can upper bound the number of rounds by analyzing how often the interval $[\log \ell, \log u]$ gets halved. This is $[\log \ell, \log u] \subseteq [\log(2^{2^a}), \log(2^{2^{a+3}})] = [2^a, 8 * 2^a]$ can be halved at most a constant number of times, until it contains only one value, which implies that $4 \cdot \ell > u$ holds.

Lemma 4.3. For a given $\mathcal{F}(t)$ and a given interval $L = [\ell, u]$ for which (P3) holds, there is an algorithm \mathcal{A}_3 that witnesses the correctness of $\mathcal{F}(t)$ until a time t' at which it outputs $L' = [\ell', u']$ for which (P3) does not hold. The algorithm uses $\mathcal{O}(\log 1/\varepsilon)$ messages on expectation.

Proof. The algorithm applies the generic framework and uses the midpoint strategy starting with the interval $L_0 := [\ell, u]$. Observe that it takes at most $\mathcal{O}(\log \frac{1}{\varepsilon})$ redefinitions of L to have the final size, no matter whether the algorithm observes only filter-violations from nodes $i \in \mathcal{F}(t)$ or $i \notin \mathcal{F}(t)$. This together with the use of the EXISTENCEPRO-TOCOL for handling filter-violations yields the needed number of messages on expectation. The correctness follows similarly as shown for Lemma 4.1.

Now we propose an algorithm started at a time t which computes the output $\mathcal{F}(t)$ and witnesses its correctness until some (not predefined) time t' at which the TOP-K-PROTOCOL terminates using a combination of the algorithms stated above. Precisely the TOP-K-PROTOCOL is defined as follows:

Algorithm TOP-K-PROTOCOL

- 1. Compute the nodes holding the (k + 1) largest values and define $\ell \coloneqq v_{k+1}^t$, $u \coloneqq v_k^t$ and $\mathcal{F}(t)$.
- 2. If (P1) holds, call \mathcal{A}_1 with the arguments $\mathcal{F}(t)$ and $L = [\ell, u]$. At the time t' at which \mathcal{A}_1 outputs $L' = [\ell', u']$ set $\ell := \ell'$ and u := u'.
- 3. If (P2) holds, call A_2 with the arguments $\mathcal{F}(t)$ and $L = [\ell, u]$. At the time t' at which A_2 outputs $L' = [\ell', u']$ set $\ell := \ell'$ and u := u'.
- 4. If (P3) holds, call A_3 with the arguments $\mathcal{F}(t)$ and $L = [\ell, u]$. At the time t' at which A_3 outputs $L' = [\ell', u']$ set $\ell := \ell'$ and u := u'.

- 5. If $u \ge \ell$ and $u \le \frac{1}{(1-\varepsilon)}\ell$ holds, set the filters to $F_1 \coloneqq [\ell, \infty]$, $F_2 \coloneqq [0, u]$. At the time t' at which node $i \in \mathcal{F}_2$ reports a filter-violation from below define $\ell \coloneqq v_i^{t'}$. In case node $i \in \mathcal{F}_1$ reports a filter-violation from above, define $u \coloneqq v_i^{t'}$.
- 6. Terminate and output (ℓ, u) .

Lemma 4.4. Consider a time t. The algorithm TOP-K-PROTOCOL computes the top-k set and witnesses its correctness until a time t' at which it outputs $L = [\ell, u]$, where $\ell \leq MAX_{\mathcal{F}_2}(t, t')$, $MIN_{\mathcal{F}_1}(t, t') \leq u$, and $\ell > u$ holds (i.e. L is empty). The algorithm uses $\mathcal{O}(k \log n + \log \log \Delta + \log \frac{1}{\epsilon})$ messages on expectation.

Proof. We first argue on the correctness of TOP-K-PROTOCOL and afterwards shortly analyze the number of messages used.

The algorithm computes in step 1. a correct output \mathcal{F}_1 at time t by using the algorithm from Lemma 2.6 for k times. In consecutive time steps t' > t the correctness of TOP-K-PROTOCOL follows from the correctness of algorithms $\mathcal{A}_1, \mathcal{A}_2$, and \mathcal{A}_3 in steps 2. - 4. For the correctness of step 5. observe that by setting the filters to $F_1 = [\ell, \infty]$ and $F_2 = [0, u]$ and the fact that $u \leq \frac{1}{1-\varepsilon}\ell$ holds the filters are valid. Thus, as long as all nodes observe values which are inside their respective filters the output need not change.

At the time step t' the protocol terminates and outputs $L = [\ell, u]$ it holds $u < \ell$. Thus, there are nodes $i_1 \in \mathcal{F}_1$ and $i_2 \in \mathcal{F}_2$ and time steps $t_1, t_2 \in [t, t']$ with: $v_{i_1}^{t_1} \leq u$ and $v_{i_2}^{t_2} \geq \ell$, and thus, $v_{i_1}^{t_1} < v_{i_2}^{t_2}$. To argue on the number of messages observe that the first step can be executed using $\mathcal{O}(k \log n)$ number of

To argue on the number of messages observe that the first step can be executed using $\mathcal{O}(k \log n)$ number of messages. At the time the condition of steps 2. - 5. are checked these steps can be performed using $\mathcal{O}(k \log n)$ number of messages, by computing the nodes holding the k + 1 largest values. The algorithms $\mathcal{A}_1, \mathcal{A}_2$, and \mathcal{A}_3 are called at most once each thus the conditions are also checked at most once. After executing step 5. the algorithm terminates which leads to the result on the number of messages as stated above.

Theorem 4.5. The algorithm TOP-K-PROTOCOL has a competitiveness of $\mathcal{O}(k \log n + \log \log \Delta + \log \frac{1}{\varepsilon})$ allowing an error of ε compared to an optimal offline algorithm that solves the exact Top-k-Position Monitoring problem.

Proof. The correctness of TOP-K-PROTOCOL and the number of messages follow from Lemma 4.4. Now we argue that OPT had to communicate at least once in the interval [t, t'] during which TOP-K-PROTOCOL was applied. If OPT communicated, the bound on the competitiveness directly follows. Now assume that OPT did not communicate in the interval [t, t']. We claim that the interval L maintained during TOP-K-PROTOCOL always satisfies the invariant $L^* \subseteq L$. If this claim is true, we directly obtain a contradiction to the fact that OPT did not communicate because of the following reasons. On the one hand, because OPT has to monitor the exact Top-k-Positions, OPT chooses the same set of nodes $\mathcal{F}^* = \mathcal{F}_1$ which was chosen by the online algorithm. On the other hand, at the time t' the algorithm TOP-K-PROTOCOL terminates, $u' < \ell'$ holds. Thus, the interval L' is empty and since $L^* \subseteq L'$ holds, it follows that L^* is empty and hence, OPT must have communicated.

We now prove the claim. Recall that TOP-K-PROTOCOL is started with an interval L that fulfills $L^* \subseteq L$ by definition. To show that $L^* \subseteq L$ holds during the entire interval [t, t'], it suffices to argue that each of the previous algorithms makes sure that when started with an interval L such that $L^* \subseteq L$, it outputs L' with $L^* \subseteq L'$. Our following reasoning is generic and can be applied to the previous algorithms. Consider the cases in which filterviolations are observed and hence the interval L is modified: If a filter-violation from below happened at a time $t_1 > t$, there is a node $i \in \mathcal{F}_2$ with a value $v_i^{t_1} > \ell'$ and thus, $\ell^* > \ell'$ holds. If a filter-violation from above happened at a time t', there is a node $i \in \mathcal{F}_1$ with a value $v_i^{t'} < u'$ and thus, $u^* < u'$ holds. This case-distinction leads to the result, that L^* has to be a subset of $[\ell', u']$.

5 Competing against an Approximate Adversary

In this section, we study the case in which the adversary is allowed to use an approximate filter-based offline algorithm, i.e. one that solves ε -Top-*k*-Position Monitoring. Not surprisingly, it turns out that it is much more challenging for online than for offline algorithms to cope with or exploit the allowed error in the output. This fact is formalized in the lower bound in Theorem 5.1, which is larger than previous upper bounds for the exact problem. However, we also propose two online algorithms that are competitive against offline algorithms that are allowed to have the same error ε and a smaller error $\varepsilon' \leq \frac{\varepsilon}{2}$, respectively.

5.1 Lower Bound for Competitive Algorithms

We show a lower bound on the competitiveness proving any online algorithm has to communicate at least $(\sigma - k)$ times in contrast to an offline algorithm which only uses k + 1 messages. Recall that the adversary generates the data streams and can see the filters communicated by the server. Note that as long as the online and the offline algorithm are allowed to make use of an error $\varepsilon \in (0, 1)$ the lower bound holds, even if the errors are different.

Theorem 5.1. Any filter-based online algorithm which solves the ε -Top-k-Position Monitoring problem and is allowed to make use of an error of $\varepsilon \in (0, 1)$ has a competitiveness of $\Omega(\sigma/k)$ compared to an optimal offline algorithm which is allowed to use a (potentially different) error of $\varepsilon' \in (0, 1)$.

Proof. Consider an instance in which the observed values of $\sigma \in [k + 1, n]$ nodes are equal to some value y_0 (the remaining $n - \sigma$ nodes observe smaller values) at time t = 0 and the following adversary: In time step $r = 0, 1, \ldots, n - k$, the adversary decides to change the value of one node i with $v_i^r = y_0$ to be $v_i^{r+1} = y_1 < (1 - \varepsilon) \cdot y_0$ such that a filter-violation occurs. Observe that such a value y_1 exists if $\varepsilon < 1$ holds and a node i always exists since otherwise the filters assigned by the online algorithm cannot be feasible. Hence, the number of messages sent by the online algorithm until time step n - k is at least n - k. In contrast, the offline algorithm knows the n - k nodes whose values change over time and hence, can set the filters such that no filter-violation happens. The offline algorithm sets two different filters: One filter $F_1 = [y_0, \infty]$ for those k nodes which have a value of y_0 at time step n - k using k messages and one filter $F_2 = [0, y_0]$ for the remaining n - k nodes using one broadcast message. By essentially repeating these ideas, the input stream can be extended to an arbitrary length, obtaining the lower bound as stated. \Box

5.2 Upper Bounds for Competitive Algorithms

Now we propose an algorithm DENSEPROTOCOL and analyze the competitiveness against an optimal offline algorithm in the setting that both algorithms are allowed to use an error of ε .

The algorithm DENSEPROTOCOL is started a time t. For sake of simplicity we assume that the k-th and the (k + 1)-st node observe the same value z, that is $z \coloneqq v_{\pi(k,t)}^t = v_{\pi(k+1,t)}^t$. However, if this does not hold we can define the filters to be $F_1 = [v_{\pi(k+1,t)}^t, \infty]$ and $F_2 = [0, v_{\pi(k,t)}^t]$ until a filter-violation is observed at some time t' using $\mathcal{O}(k \log n)$ messages on expectation. If the filter-violation occurred from below define $z \coloneqq v_{\pi(k,t)}^t$ and if a filter-violation from above is observed define $z \coloneqq v_{\pi(k+1,t)}^t$.

filter-violation from above is observed define $z := v_{\pi(k+1,t)}^t$. The high-level idea of DENSEPROTOCOL is similar to the TOP-K-PROTOCOL to compute a guess L on the lower endpoint of the filter of the output \mathcal{F}^* of OPT (assuming OPT did not communicate during [t, t']) for which the invariant $\ell^* \in L^* \subseteq L_r$ holds. The goal of DENSEPROTOCOL is to halve the interval L while maintaining $\ell^* \in L$ until $L = \emptyset$ and thus show that no value exists which could be used by OPT.

To this end, the algorithm partitions the nodes into three sets. Intuitively speaking, the first set which we call V_1 contains those nodes which have to be part of the optimal output, V_3 those nodes that cannot be part of any optimal output and V_2 the remaining nodes. The sets change over time as follows. Initially V_1^t contains those nodes that observes a value $v_i^t > \frac{1}{1-\varepsilon}z$. Since the algorithm may discover at a time t' > t that some node i has to be moved to $V_1^{t'+1}$ which also contains all nodes from previous rounds, i.e. $V_1^{t'} \subseteq V_1^{t'+1}$. On the other hand V_3^t initially contains the nodes which observed a value $v_i^t < (1-\varepsilon)z$. Here also the algorithm may discover at a time t' > t that some node i has to be moved to $V_3^{t'+1}$ which (similar to V_1) contains nodes from previous rounds. At the time t the set V_2^t simply contains the remaining nodes $\{1, \ldots, n\} \setminus (V_1^t \cup V_3^t)$ and its cardinality will only decrease over time.

In the following we make use of sets S_1 and S_2 to indicate that nodes in V_2 may be moved to V_1 or V_3 depending on the values observed by the remaining nodes in V_2 . Nodes in S_1 observed a value larger than z but still not that large to decide to move it to V_1 and similarly nodes in S_2 observed smaller values than z but not that small to move it to V_3 .

Next we propose the algorithm DENSEPROTOCOL in which we make use of an algorithm SUBPROTOCOL for the scenario in which some node i exists that is in S_1 and in S_2 . At a time at which the SUBPROTOCOL terminates it outputs that ℓ^* has to be in the lower half of L or in the upper half of L thus, the interval L gets halved (which initiates the next round) or moves one node from V_2 to V_1 or V_3 . Intuitively speaking SUBPROTOCOL is designed such that, if OPT did not communicate during [t, t'], where t is the time the DENSEPROTOCOL is started and t' is the current time step, the movement of one node $i \in V_2$ to V_1 or V_3 implies that i has necessarily to be part of \mathcal{F}^* or not. For now we assume the algorithm SUBPROTOCOL to work correctly as a black box using SUB(n, |L|) number of messages.

Note that in case L_r contains one value and gets halved, the interval L_{r+1} is defined to be empty. In case the algorithm observes multiple nodes reporting a filter-violation the server processes one violation at a time in an arbitrary

order. Since the server may define new filters after processing a violation one of the multiple filter-violations may be not relevant any longer, thus the server simply ignores it.

Algorithm: DENSEPROTOCOL

1. Define $z \coloneqq v_k^t = v_{k+1}^t$ and the following sets: $V_1 \coloneqq \{i \in \{1, ..., n\} \mid v_i^t > \frac{1}{1-\varepsilon}z\},\$ $V_3 \coloneqq \{i \in \{1, ..., n\} \mid v_i^t < (1-\varepsilon)z\},\$ $V_2 \coloneqq \{1, ..., n\} \setminus (V_1 \cup V_3).$

Define an interval $L_0 := [(1 - \varepsilon)z, z]$ and define sets S_1, S_2 of nodes which are initially empty and use S to denote $S_1 \cup S_2$. Set r := 0 indicating the round.

- 2. The following **rules** are applied for (some) round r: Let ℓ_r be the midpoint of L_r and $u_r := \frac{1}{1-\varepsilon}\ell_r$ For a node i the filter is defined as follows: If $i \in V_1, F_i := [\ell_r, \infty]$; If $i \in V_2 \cap S_1, F_i := [\ell_r, \frac{1}{1-\varepsilon}z]$. if $i \in V_2 \setminus S, F_i := [\ell_r, u_r]$; If $i \in V_2 \cap S_2, F_i := [(1-\varepsilon)z, u_r]$. if $i \in V_3, F_i := [0, u_r]$. The output $\mathcal{F}(t)$ is defined as $V_1 \cup (S_1 \setminus S_2)$ and $k - |V_1 \cup (S_1 \setminus S_2)|$ many nodes from $V_2 \setminus S_2$.
- 3. Wait until time t', at which some node *i* reports a filter-violation:
 - a. If $i \in V_1$, then set L_{r+1} to be the lower half of L_r and define $S_2 \coloneqq \emptyset$.
 - b. If $i \in (V_2 \setminus S)$ violates its filter from below then
 - b.1. If the server observed strictly more than k nodes with larger values than u_r then set L_{r+1} to be the upper half of L_r and define $S_1 := \emptyset$.
 - b.2. else add i to S_1 and update i's filter.
 - c. If $i \in S_1 \setminus S_2$ violates its filter then
 - c.1. If i violates its filter from below then move i from S_1 and V_2 to V_1 and update i's filter.
 - c.2. else add i to S_2 and call SUBPROTOCOL.
 - d. If the server observed k nodes with values $v_i > u_r$ and n k nodes with values $v_i < \ell_r$ then call TOP-K-PROTOCOL
 - e. If L_{r+1} was set if is empty, end the protocol, otherwise increment r, update u_r , ℓ_r , all filters using the rules in 2., and go to step 3.
 - And their symmetric cases -
 - a'. If $i \in V_3$ then set L_{r+1} to be the upper half of L_r and define $S_1 := \emptyset$.
 - b'. If $i \in (V_2 \setminus S)$ violates its filter from above then
 - b'.1. If the server observed strictly more than n k nodes with smaller values than ℓ_r then set L_{r+1} to the lower half of L_r and define $S_2 := \emptyset$.
 - b'.2. else add i to S_2 .
 - c'. If $i \in S_2 \setminus S_1$ violates its filter then
 - c'.1. If *i* violates its filter from above then delete *i* from S_2 , delete *i* from V_2 , and add *i* to V_3 .
 - c'.2. else add i to S_1 and call SUBPROTOCOL.

We analyze the correctness of the protocol in the following lemma and the number of messages used in Lemma 5.3. We prove that OPT communicated at least once in Lemma 5.7.

Lemma 5.2. The protocol DENSEPROTOCOL computes a correct output $\mathcal{F}(t')$ at any time t'.

Proof. By definition the output consists of nodes from V_1 , S_1 and (arbitrary) nodes from $V_2 \setminus S_2$ (cf. step 2.). Observe that by definition of the filters of the nodes in these subsets, the minimum of all lower endpoints of the filters is ℓ_r following the rules in step 2. Also observe that the maximum of all upper endpoints of the filters of the remaining nodes is u_r . Since by definition $u_r = \frac{1}{1-\varepsilon}\ell_r$ holds, the values observed by nodes $i \in \mathcal{F}_1$ are (lower) bounded by ℓ_r and nodes $i \in \mathcal{F}_2$ are (upper) bounded by u_r , thus the overlap of the filters is valid.

Now we argue that there are at least k nodes in the set $V_1 \cup S_1 \cup V_2 \setminus S_2$. To this end, assume to the contrary that t' is the first time step at which strictly less than k nodes are in the union of these sets. Now observe that the cases in the DENSEPROTOCOL in which nodes are deleted from one of V_1, S_1 or $V_2 \setminus S_2$ are 3.c.1., 3.c.2., and 3.b'.2..

Observe that in step 3.c.1. the algorithm moves i from S_1 and V_2 to V_1 and thus i is again part of the output and does not change the cardinality. In step 3.c.2. the node i is added to S_2 and SUBPROTOCOL is called afterwards. At this time t' node i is (again) part of the output of SUBPROTOCOL and thus there are sufficiently many nodes to choose as an output which is a contradiction to the assumption. In the remaining case 3.b'.2. DENSEPROTOCOL adds i to S_2 . However, since at time t' strictly less than k nodes are in $V_1 \cup S_1 \cup (V_2 \setminus S_2)$, there are strictly more than n - k nodes in $S_2 \cup V_3$ and thus, the algorithm would execute step 3.b'.1. instead. This leads to a contradiction to the assumption. By these arguments the correctness follows.

Lemma 5.3. The protocol DENSEPROTOCOL uses at most $O(k \log n + \sigma \log(\varepsilon v_k) + (\sigma + \log(\varepsilon v_k)) \cdot SUB(\sigma, |L|))$ messages on expectation.

Proof. Initially the algorithm computes the top-k set and probes all nodes which are in the ε -neighborhood of the node observing the k-th largest value, using $O(k \log n + \sigma)$ messages on expectation.

During each round r each node can only violate its filter at most constant times without starting the next round r + 1 or leading to a call of SUBPROTOCOL based on the following simple arguments: All nodes i in V_1 or V_3 directly start the next round r + 1 after a filter-violation. Now fix a node $i \in V_2$ and observe that if it is not contained in S_1 and S_2 it is added to S_1 if a filter-violation from below or to S_2 if a filter-violation from above is observed. At the time this node i observes a filter-violation in the same direction (i.e. from below if it is in S_1 and from above if it is in S_2) it is added to V_1 or V_3 . In these cases the next filter-violation will start the next round. The last case that remains is that it is added to both sets, S_1 and S_2 . Observe that the SUBPROTOCOL is called and starts the next round or decides on one node (which may be different from the fixed node i) to be moved to V_1 or V_3 .

Observe that at most $\sigma + 1$ nodes can perform filter-violations without starting the next round since each node from V_1 or V_3 directly starts the next round and the number of nodes in V_2 is bounded by σ . Furthermore observe that after each round the interval L is halved thus, after at most $\log |L_0| + 1$ rounds the set L_r is empty.

Now focus on the SUBPROTOCOL which also halves L after termination or decides on one node $i \in V_2^{t'}$ to be moved to $V_1^{t'+1}$ or $V_3^{t'+1}$. Thus, it can be called at most $\sigma + \log(\varepsilon v_k)$ times, leading to the result as stated above. \Box

The SUBPROTOCOL We propose an algorithm which is dedicated for the case in the execution of DENSEPROTO-COL that one node i was added to S_1 and to S_2 .

 $\in V_2 \setminus S$ reported a filter-violation from below and from above and thus gets added to S_1 and to S_2 (in an arbitrary order). In detail, it has observed a value which is larger than u_r and a value which is smaller than ℓ_r . As a short remark, if $i \in \mathcal{F}^*$ would hold, then $\ell^* \leq \ell_r$ follows and on the other hand if $i \notin \mathcal{F}^*$ holds, then $\ell^* \geq \ell_r$ follows, but in DENSEPROTOCOL cannot decide $i \in \mathcal{F}^*$ in steps 3.c.2. or 3.c.2.

Algorithm: SUBPROTOCOL

- 1. Define an interval $L'_0 \coloneqq L_r \cap [(1 \varepsilon)z, \ell_r], S'_1 \coloneqq S_1$, and $S'_2 \coloneqq \emptyset$. Set $r' \coloneqq 0$ indicating the round.
- 2. The following **rules** are applied for (some) round r': Let ℓ'_r be the midpoint of $L'_{r'}$ and $u'_{r'} \coloneqq \frac{1}{1-\varepsilon}\ell'_{r'}$. For a node i the filter is defined as follows: If $i \in V_1$, $F'_i \coloneqq F_i$; If $i \in V_2 \cap (S'_1 \setminus S'_2)$, $F'_i \coloneqq [\ell_r, \frac{1}{1-\varepsilon}z]$. If $i \in V_2 \cap S'_1 \cap S'_2$, $F'_i \coloneqq [\ell'_{r'}, \frac{1}{1-\varepsilon}z]$; if $i \in V_2 \cap (S'_1 \setminus S'_2)$, $F'_i \coloneqq [\ell_r, u'_{r'}]$. if $i \in V_2 \cap (S'_2 \setminus S'_1)$, $F'_i \coloneqq [(1-\varepsilon)z, u'_{r'}]$; if $i \in V_3$, $F'_i \coloneqq [0, u'_{r'}]$;

The output $\mathcal{F}(t)$ is defined as $V_1 \cup (S'_1 \setminus S'_2) \cup (S'_1 \cap S'_2)$ and sufficiently many nodes from $V_2 \setminus S'_2$.

- 3. Wait until time t', at which node *i* reports a filter-violation:
 - a. If $i \in V_1$, then terminate SUBPROTOCOL and set L_{r+1} to be the lower half of L_r .
 - b. If $i \in (V_2 \setminus S')$ violates its filter from below
 - b.1. If the server observed strictly more than k nodes with larger values than u_r then
 - set $L'_{r'+1}$ to be the upper half of $L'_{r'}$ and redefine $S'_1 := S_1$.
 - If $L'_{r'+1}$ is defined to the empty set **then** terminate SUBPROTOCOL and define the last node i which was in $S'_1 \cap S'_2$ and observed a filter-violation from above to be moved to V_3 . If such a node does not exist the node $i \in S_1 \cap S_2$ moves to V_3 .
 - b.2. Else add i to S'_1 .
 - c. If $i \in S'_1 \setminus S'_2$ violates its filter
 - c.1. If *i* violates its filter from below then move *i* from V_2 and S'_1 to V_1 .
 - c.2. Else add i to S'_2 and update i'th filter.
 - d. If $i \in S'_1 \cap S'_2$ violates its filter
 - d.1. If i violates from below then move i to V_1 terminate the SUBPROTOCOL.
 - d.2. else
 - define $L'_{r'+1}$ to be the lower half of $L'_{r'}$ and redefine $S'_2 \coloneqq \emptyset$.
 - If $L'_{r'+1}$ is defined to be the empty set **then** terminate SUBPROTOCOL and move *i* to V_3 .
 - e. If the server observed k nodes with values $v_i > u_r$ and n k nodes with values $v_i < \ell_r$ then call TOP-K-PROTOCOL
 - f. If $L'_{r'+1}$ was set increment r', update $u'_{r'}$, $\ell'_{r'}$, all filters using the rules in 2., and go to step 3.

- And their symmetric cases -

- a'. If $i \in V_3$, then
 - set $L'_{r'+1}$ to be the upper half of $L'_{r'}$ and redefine $S'_1 \coloneqq S_1$.
 - If $L'_{r'+1}$ is defined to the empty set then terminate SUBPROTOCOL and define the last node i which was in $S'_1 \cap S'_2$ and observed a filter-violation from above to be moved to V_3 . If such a node does not exist the node $i \in S_1 \cap S_2$ moves to V_3 .
- b'. If $i \in (V_2 \setminus S')$ violates its filter from above
 - b'.1. If the server observed strictly more than n k nodes with a value less than ℓ_r , then terminate SUB-PROTOCOL and set L_{r+1} to be the lower half of L_r .
 - b'.2. else add i to S'_2 .
- c'. If $i \in S'_2 \setminus S'_1$
 - c'.1. If *i* violates its filter from above then move *i* from V_2 and S'_2 to V_3 .
 - c'.2. else add i to S'_1 and update i'th filter.

Lemma 5.4. The protocol SUBPROTOCOL computes a correct output $\mathcal{F}(t')$ at any time t' at which a node $i \in S_1 \cap S_2$ exists.

Proof. By definition the output consists of nodes from V_1 , $S'_1 \setminus S'_2$, $S'_1 \cap S'_2$ and (arbitrary) nodes from $V_2 \setminus S'_2$ (cf. step 2.). Observe that by definition of the filters of the nodes in these subsets, the minimum of all lower endpoints of the filters is $\ell'_{r'}$ (in case the node is in S_1 and in S_2) following the rules in step 2. Also observe that the maximum of all upper endpoints of the filters of the remaining nodes (in subsets $V_2 \setminus S'$, $S'_2 \setminus S'_1$ or V_3) is $u'_{r'}$. Since by definition $u'_{r'} = \frac{1}{1-\varepsilon}\ell'_{r'}$ holds, the values observed by nodes $i \in \mathcal{F}_1$ are (lower) bounded by $\ell'_{r'}$ and nodes $i \in \mathcal{F}_2$ are (upper) bounded by $u'_{r'}$ thus, the overlap of the filters is valid.

Now we argue that there are at least k nodes in the sets V_1 , $S_1 \setminus S_2$, $S_1 \cap S_2$, and $V_2 \setminus S_2$. To this end, simply assume to the contrary that at a time t' there are strictly less than k nodes in the union of the sets. It follows that at

this time t', the algorithm has observed that there are strictly more than n - k nodes with a value smaller than $\ell'_{r'}$. Thus, the algorithm would continue (compare case b'.1.) with a lower value of ℓ_r or, in case the interval L_r is empty, terminates (which is a contradiction).

By these arguments the correctness follows.

Lemma 5.5. The protocol SUBPROTOCOL uses at most $\mathcal{O}(\sigma \log |L|)$ messages on expectation.

Proof. During each round r' each node can only violate its filter at most constant times without starting the next round r' + 1 based on the following simple arguments: All nodes i in V_1 or V_3 directly start the next round r' + 1 after a filter-violation. Now fix a node $i \in V_2$ and observe that if it is not contained in S'_1 and S'_2 it is added to S'_1 if a filter-violation from below or to S'_2 if a filter-violation from above is observed. At the time this node i observes a filter-violation in the same direction (i.e. from below if it is in S'_1 and from above if it is in S'_2) it is added to V_1 or V_3 . In these cases the next filter-violation will start the next round. The last case that remains is that it is added to both sets, S'_1 and S'_2 . Observe that TOP-K-PROTOCOL terminates if $i \in S'_1 \cap S'_2$ violates its filter from below (and moves $itoV_1$). Otherwise i violates its filter from above SUBPROTOCOL starts the next round r' + 1.

Observe that at most $\sigma + 1$ nodes can perform filter-violations without starting the next round since each node from V_1 or V_3 directly starts the next round (r + 1 from the DENSEPROTOCOL or r' + 1 this protocol) and the number of nodes in V_2 is bounded by σ .

Furthermore observe that after each round the interval L', the guess of OPTs lower endpoint of the upper filter, is halved. The range of L' is upper bounded by the range of L thus, after at most $\log |L| + 1$ rounds the set L' is empty.

Lemma 5.6. Given a time point t at which SUBPROTOCOL is started. At the time t' which SUBPROTOCOL terminates, there is one node i that is moved from V_2 to V_1 or V_3 or the interval L_r (from DENSEPROTOCOL) is halved correctly.

Proof. Focus on the cases in which L' is halved or there is a decision on a node i to move to V_1 or V_3 (cf. cases 3.b.1., 3.d.1. 3.d.2., 3.a'., and 3.c'.1.).

In step 3.b.1. the server observed at the time t' a filter-violation from $i \in V_2 \setminus S'$ and there are (strictly) more than k nodes observed with a larger value than $u'_{r'}$. Observe that in this case for all subsets S with k elements there exists one node $i \notin S$ which observed a value $v_i \ge u'_{r'}$, thus no matter which set is chosen by OPT, for the upper bound u^* for nodes $i \notin \mathcal{F}^*$ it holds: $u^* \ge u'_{r'}$, and since $u'_{r'} = \frac{1}{1-\varepsilon}\ell'_{r'}$ holds, it follows $\ell^* \ge \ell'_{r'}$. Furthermore if $L'_{r'+1}$ was defined as the empty set, and a node $i \in S'_1 \cap S'_2$ exists, observe that i gets a value $v_i \le \ell'_{r'}$ and since in this case $u^* \ge u'_{r'}$ holds, $i \notin \mathcal{F}^*$ follows. If such a node i does not exist during the execution of SUBPROTOCOL, the node $i \in S_1 \cap S_2$ which initiated the SUBPROTOCOL can be decided to move to V_3 since during the execution of SUBPROTOCOL the interval L' is only halved to the upper half, thus $i \in S_1 \cap S_2$ observed a value $v_i < \ell_r = \ell'_{r'}$ and since $u^* \ge u'_{r'}$ holds, this $i \notin \mathcal{F}^*$ follows.

In step 3.d.1. the node *i* observed a value v_i which is larger than $\frac{1}{1-\varepsilon}z$ and thus has to be part of \mathcal{F}^* .

In step 3.d.2. the node *i* observed a value $v_i < \ell'_{r'}$. If during the execution of SUBPROTOCOL the set L' was defined as the upper half at least once then there was a node $j \in V_3$ or strictly more than *k* nodes which observed a larger value than $u'_{r'}$. It follows, that this *i* cannot be part of \mathcal{F}^* . In case during the execution of SUBPROTOCOL the set L' is alway defined to the lower half, then $\ell'_{r'}$ is the lower end of *L* and since node *i* observed a value strictly smaller than $\ell'_{r'}$ it cannot be part of \mathcal{F}^* .

The arguments for case 3.a'. are similar to 3.b.1.

For the remaining case 3.c'.1. simply observe that *i* observed a smaller value than $(1 - \varepsilon)z$ thus *i* cannot be part of \mathcal{F}^* follows.

First, focus on the steps in which L is halved and observe that steps 3.a. and 3.b'.1. are the same cases as in the DENSEPROTOCOL.

Lemma 5.7. Given a time point t at which DENSEPROTOCOL is started. Let t' be the time point at which DENSE-PROTOCOL terminates. During the time interval [t, t'] OPT communicated at least once.

Proof. We prove that OPT communicated by arguing that ℓ^* , the lower endpoint of the upper filter, i.e. the filter for the output \mathcal{F}^* , is in the guess L_r at each round r ($\ell^* \in L^* \subseteq L_r$). Hence we show that although if we halve the interval L_r , the invariant $\ell^* \in L^* \subseteq L_r$ is maintained all the time of the execution of DENSEPROTOCOL and possible calls of SUBPROTOCOL.

In the following we assume to the contrary that OPT did not communicate throughout the interval [t, t']. We first argue for the execution of DENSEPROTOCOL and assume that the invariant by calls of SUBPROTOCOL hold by Lemma 5.6.

First focus on the DENSEPROTOCOL, which halves the interval L_r in steps 3.a., 3.b.1., 3.a'., and 3.b'.1.:

In step 3.a. in which a node $i \in V_1$ violates its filter from above and observes a value $v_i < \ell_r$, it holds: $i \in \mathcal{F}^*$ thus, $\ell^* < \ell_r$ follows.

In step 3.b.1. there are (strictly) more than k nodes with a larger value than u_r . It follows that for all subsets S (with k elements) there is one node $i \notin S$ observing a value larger than u_r and thus, $\ell^* \ge (1 - \varepsilon)u_r = \ell_r$ holds.

The case 3.a'. (which is symmetric to 3.a.) is executed if a node $i \in V_3$ observed a filter-violation $(v_i > u_r)$ which implies that the upper endpoint u^* of filter F_2 is larger than v_i and thus, $\ell^* \ge (1 - \varepsilon)u_r = \ell_r$.

In step 3.b'.1. (which is symmetric to 3.b.1.) there are (strictly) more than n - k nodes with a smaller value than ℓ_r . It follows that for all subsets S (with k elements) there is one node $i \in S$ observing a value smaller than ℓ_r and thus, $\ell^* \leq \ell_r$ holds.

Theorem 5.8. There is an online algorithm for ε -Top-k-Position Monitoring which is $\mathcal{O}(\sigma^2 \log(\varepsilon v_k) + \sigma \log^2(\varepsilon v_k) + \log \log \Delta + \log \frac{1}{\varepsilon})$ -competitive against an optimal offline algorithm which may use an error of ε .

Proof. The algorithm works as follows. At time t at which the algorithm is started, the algorithm probes the nodes holding the k + 1 largest values. If $v_{\pi(k+1,t)}^t < (1 - \varepsilon)v_{\pi(k,t)}^t$ holds, the algorithm TOP-K-PROTOCOL is called. Otherwise the algorithm DENSEPROTOCOL is executed. After termination of the respective call, the procedure starts over again.

Observe that if the condition holds, there is only one unique output and thus, the TOP-K-PROTOCOL monitors the Top-k-Positions satisfying the bound on the competitiveness as stated in Theorem 4.5. If the condition does not hold, there is at least one value in the ε -neighborhood of $v_{\pi(k,t)}^t$ and thus, the DENSEPROTOCOL monitors the approximated Top-k-Positions as analyzed in this section.

The number of messages used is simply obtained by adding the number of messages used by the respective algorithms as stated above. $\hfill \Box$

To obtain the upper bounds stated at the beginning, we upper bound σ by n and v_k by Δ : $\mathcal{O}(n^2 \log(\varepsilon \Delta) + n \log^2(\varepsilon \Delta) + \log \log \Delta + \log \frac{1}{\varepsilon})$. Note that for constant ε we obtain a slightly simpler bound of $\mathcal{O}(n^2 \log \Delta + n \log^2 \Delta)$ on the competitiveness.

Corollary 5.9. There is an online algorithm for ε -Top-k-Position Monitoring which is $\mathcal{O}(\sigma + k \log n + \log \log \Delta + \log \frac{1}{\varepsilon})$ -competitive against an optimal offline algorithm which may use an error of $\varepsilon' \leq \frac{\varepsilon}{2}$.

Proof. The algorithm works as follows. At the initial time step t the algorithm probes the nodes holding the k + 1 largest values. If $v_{\pi(k+1,t)}^t < (1 - \varepsilon)v_{\pi(k,t)}^t$ holds the algorithm TOP-K-PROTOCOL is called. Otherwise the online algorithm simulates the first round of the DENSEPROTOCOL, that is nodes are partitioned

Otherwise the online algorithm simulates the first round of the DENSEPROTOCOL, that is nodes are partitioned into V_1, V_2 , and V_3 and the filters are defined as proposed (cf step 2. of DENSEPROTOCOL). Here all nodes with values larger than $\frac{1}{1-\varepsilon}(1-\frac{\varepsilon}{2})z$ are directly added to V_1 instead of adding to S_1 , and nodes observing values smaller than $(1-\frac{\varepsilon}{2})z$ are added to V_3 . Furthermore, if a filter-violation from some node $i \in V_2$ is observed, it is directly moved (deleted from V_2 and added) to V_1 in case it violates from below, and added to V_3 if violated from above.

Whenever a node from V_1 (or from V_3) violates its filter the algorithm terminates. Additionally if (strictly) more than k nodes are in V_1 the algorithm is terminated or if (strictly) less than k nodes are in $V_1 \cup V_2$. If exactly k nodes are in V_1 and n - k nodes are in V_3 the TOP-K-PROTOCOL is executed.

For the following argumentation on the competitiveness we focus on the case that TOP-K-PROTOCOL was not called since the analysis of TOP-K-PROTOCOL holds here. Observe that OPT (with an error of ε') had to communicate based on the following observation:

Let t' be the time at which the algorithm terminates. Assume to the contrary that OPT did not communicate during [t, t']. In case node $i \in V_1$ observes a filter-violation from above, $(v_i < (1 - \frac{\varepsilon}{2})z)$ and $\varepsilon' \leq \frac{\varepsilon}{2}$, OPT had to set $\ell^* \leq v_i$ and $u^* \geq z$, which leads to a contradiction to the definition of filters. In case node $i \in V_3$ observes a filter-violation from below, $(v_i > \frac{1}{1-\varepsilon}(1-\frac{\varepsilon}{2})z)$ and $\varepsilon' \leq \frac{\varepsilon}{2}$, OPT had to set $u^* \geq v_i$ and $\ell^* \leq z$, which leads to a contradiction to the definition of filters. In case node $i \in V_3$ observes a filter-violation to the definition of filters. The fact that OPT had to communicate in the remaining cases follows by the same arguments. Since all cases lead to a contradiction, the bound on the competitiveness as stated above follows.

References

- Arackaparambil, C., Brody, J., Chakrabarti, A.: Functional Monitoring without Monotonicity. In: Proceedings of the 36th International Colloquium on Automata, Languages and Programming, pp. 95–106. Springer, Berlin (2009)
- [2] Cormode, G.: The Continuous Distributed Monitoring Model. ACM SIGMOD Record 42.1, pp. 5–14. (2013)
- [3] Cormode, G., Muthukrishnan, S., Ke, Y.: Algorithms for Distributed Functional Monitoring. ACM Transactions on Algorithms 7, 21 (2011)
- [4] Giannakopoulos Y., Koutsoupias, E.: Competitive Analysis of Maintaining Frequent Items of a Stream. Theoretical Computer Science 562, pp. 23–32. (2105)
- [5] Lam, T.W., Liu, C.-M., Ting, H.-F.: Online Tracking of the Dominance Relationship of Distributed Multidimensional Data. In: Proceedings of the 8th International Workshop on Approximation and Online Algorithms, pp. 178–189. Springer, (2011)
- [6] Mäcker, A., Malatyali, M., Meyer auf der Heide, F.: Online Top-k-Position Monitoring of Distributed Data Streams. In: Proceedings of the 29th International Parallel and Distributed Processing Symposium, pp. 357–364. IEEE, (2015)
- [7] Muthukrishnan, S.: Data Streams: Algorithms and Applications. Now Publishers Inc, (2005)
- [8] Phillips, J., Verbin, E., Zhang, Q.: Lower Bounds for Number-in-Hand Multiparty Communication Complexity, Made Easy. In: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 386–501. SIAM (2012)
- [9] Sanders, P., Schlag, S., Müller, I.: Communication Efficient Algorithms for Fundamental Big Data Problems. In: Proceedings of the IEEE International Conference on Big Data, pp. 15–23. IEEE, Silicon Valley (2013)
- [10] Tang M., Li F., Tao Y.: Distributed Online Tracking. In: Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, pp. 2047–2061. ACM, (2015)
- [11] Yi, K., Zhang, Q.: Multidimensional Online Tracking. ACM Transactions on Algorithms 8, 12 (2012)
- [12] Zhang, Z., Cheng, R., Papadias, D. and Tung, A.K.H.: Minimizing the Communication Cost for Continuous Skyline Maintenance. In: Proceedings of the ACM SIGMOD International Conference on Management of data, pp. 495–508. ACM, New York (2009)
- [13] Zhang, Q.: Communication-Efficient Computation on Distributed Noisy Datasets. In: Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures, pp. 313–322. ACM, (2015)