

A Geometric Approach to Monitoring Threshold Functions over Distributed Data Streams

IZCHAK SHARFMAN and ASSAF SCHUSTER

Technion

and

DANIEL KEREN

Haifa University

Monitoring data streams in a distributed system is the focus of much research in recent years. Most of the proposed schemes, however, deal with monitoring simple aggregated values, such as the frequency of appearance of items in the streams. More involved challenges, such as the important task of feature selection (e.g., by monitoring the information gain of various features), still require very high communication overhead using naive, centralized algorithms.

We present a novel geometric approach which reduces monitoring the value of a function (vis-à-vis a threshold) to a set of constraints applied locally on each of the streams. The constraints are used to locally filter out data increments that do not affect the monitoring outcome, thus avoiding unnecessary communication. As a result, our approach enables monitoring of arbitrary threshold functions over distributed data streams in an efficient manner.

We present experimental results on real-world data which demonstrate that our algorithms are highly scalable, and considerably reduce communication load in comparison to centralized algorithms.

Categories and Subject Descriptors: C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed applications*

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Distributed monitoring

ACM Reference Format:

Sharfman, I., Schuster, A., and Keren, D. 2007. A geometric approach to monitoring threshold functions over distributed data streams. *ACM Trans. Datab. Syst.* 32, 3, Article 23 (November 2007), 29 pages. DOI = 10.1145/1292609.1292613 <http://doi.acm.org/10.1145/1292609.1292613>

Partially supported by the Israeli Ministry of Science.

Authors' addresses: I. Sharfman and A. Schuster, Computer Science Department, Technion—Israel Institute of Technology, Haifa 32000, Israel; email: {tsachis,assaaf}@cs.technion.ac.il; D. Keren, Department of Computer Science, Haifa University, Haifa 31905, Israel; email: dkeren@cs.haifa.ac.il.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org. © 2007 ACM 0362-5915/2007/11-ART23 \$5.00 DOI 10.1145/1292609.1292613 <http://doi.acm.org/10.1145/1292609.1292613>

ACM Transactions on Database Systems, Vol. 32, No. 4, Article 23, Publication date: November 2007.

1. INTRODUCTION

A common requirement in many emerging applications is the ability to process, in real time, a continuous high-volume stream of data. Examples of such applications are sensor networks [Madden and Franklin 2002], real-time analysis of financial data [Yi et al. 2000; Zhu and Shasha 2002], and intrusion detection. These applications are commonly referred to as *data stream systems* [Babcock et al. 2002]. The real-time nature of data stream systems and the vast amounts of data they are required to process introduce new fundamental problems that are not addressed by traditional database management systems (DBMSs). Traditional DBMSs are based on a *pull* paradigm, where users issue queries regarding data stored by the system, and the system processes these queries as they are issued and returns results. Data stream systems [Carney et al. 2002; Cherniack et al. 2003; Liu et al. 1999; Madden and Franklin 2002; Motwani et al. 2003] are based on a *push* paradigm, where the users issue *continuous queries* [Babu and Widom 2001; Terry et al. 1992] that specify the required processing of the data, which the system processes as it arrives, continuously providing the user with updated results.

Various types of continuous queries have been studied in the past, including continuous versions of selection and join queries [Madden et al. 2002], various types of aggregation queries [Manku and Motwani 2002; Alon et al. 1996], and monitoring queries [Carney et al. 2002]. While most previous work regarding data stream systems considers sequential setups (the data is processed by a single processor), many data stream applications are inherently distributed: examples include sensor networks [Madden and Franklin 2002], network monitoring [Dilman and Raz 2001], and distributed intrusion detection [Manjhi et al. 2005].

A useful class of queries in the context of distributed data streams are monitoring queries. Previous work in the context of monitoring distributed data streams considered monitoring simple aggregates, such as detecting when the sum of a distributed set of variables exceeds a predetermined threshold [Dilman and Raz 2001], or finding frequently occurring items in a set of distributed streams [Manjhi et al. 2005]. Some work has been done on monitoring more complex constructs derived from distributed streams, but the proposed solutions are customized for the problem at hand. Examples include Bulut et al. [2005] which presented an algorithm for detecting similar sets of streams among a large set of distributed streams, and Cormode et al. [2005], which presented an algorithm for approximating quantiles over distributed streams.

A useful, more general type of monitoring query can be defined as follows: let X_1, X_2, \dots, X_d be frequency counts for d items over a set of streams. Let $f(X_1, X_2, \dots, X_d)$ be an arbitrary function over the frequency counts. We are interested in detecting when the value of $f(X_1, X_2, \dots, X_d)$ rises above, or falls below, a predetermined threshold value. We refer to this query as a *threshold function query*.

There is a fundamental difference between the cases of linear and nonlinear f , which can be demonstrated even for the case of one-dimensional data. Let x_1 and x_2 be values stored in two distinct nodes, and let $f(x) = 6x - x^2$. Suppose

one needs to determine whether $f(\frac{x_1+x_2}{2}) > 1$. If f was linear, the solution would have been simple, since in that case $f(\frac{x_1+x_2}{2}) = \frac{f(x_1)+f(x_2)}{2}$. Suppose that initially the value at each node is < 1 ; then a simple distributed algorithm for monitoring whether $f(\frac{x_1+x_2}{2}) > 1$ is for each node i to remain silent as long as $f(x_i) < 1$. However, even for the simple nonlinear function above, *it is impossible to determine from the values of f at the nodes whether its value at the average is above 1 or not*. For example, if $x_1 = 0, x_2 = 6$, then f 's value in each node is below 1, but its value in the average of x_1 and x_2 is 9. But if $x_1 = 10, x_2 = 20$, the value at both x_i and their average is below 1. So nothing can be deduced about the location of $f(\frac{x_1+x_2}{2})$ vis-a-vis the threshold given the locations of $f(x_i)$ vis-à-vis it.

In this trivial example, the cost of sending the data stored in the nodes is the same as sending the value, but in data mining applications the data can be of very high dimensionality. This necessitates a distributed algorithm for locally determining whether f 's value at the average data vector is above the threshold.

Following is a more practical example of a threshold function query: consider a classifier built over data extracted from a set of streams, for example a distributed spam mail filtering system. Such a system is comprised of agents installed on several dispersed mail servers. Users mark spam mail they have received as such, providing each server with a continuous stream of positive and negative samples. These samples serve as a basis for building a classifier. Since the vocabulary comprising these samples may be very large, an important task in such a setup is determining which words, or features, should be used for performing the classification. This task is known as *feature selection*. Feature selection is typically performed by calculating, for every feature, a nonlinear scoring function, such as information gain or χ^2 , over statistics collected from all the streams. All the features scoring above a certain threshold are chosen as parameters for the classification task. Since the characteristics of spam mail may vary over time, one may wish to monitor the features in order to determine if selected features remain prominent, or if any of the features not selected have become prominent. Determining whether a certain feature should be selected at a given time can be viewed as a threshold function query.

Threshold function queries can be implemented by collecting all the mail items to a central location, but such a solution is very costly in terms of communications load. We are interested in algorithms that implement threshold function queries in a more efficient manner. We achieve this by defining numerical constraints on the data collected at each node. As data arrives on the streams, every node verifies that the constraint on its stream has not been violated. We will show that, as long as none of these constraints have been violated, the query result is guaranteed to remain unchanged, and thus no communication is required.

In this article we present two algorithms for efficiently performing threshold function queries. The algorithms are based on a geometric analysis of the problem. Upon initialization, the algorithms collect frequency counts from all the streams, and calculate the initial result of the query. In addition, a numerical

constraint on the data received on each individual stream is defined. As data arrives on the streams, each node verifies that the constraint on its stream has not been violated. The geometric analysis of the problem guarantees that as long as the constraints on all the streams are upheld, the result of the query remains unchanged, and thus no communication is required. If a constraint on one of the streams is violated, new data is gathered from the streams, the query is reevaluated, and new constraints are set on the streams.

The first algorithm is a decentralized algorithm, designed for a closely coupled environment, where nodes can efficiently broadcast messages. The second algorithm is designed for loosely coupled environments, where the cost of broadcasting a message is high. These algorithms are, to the best of our knowledge, the first to enable efficient monitoring of arbitrary threshold functions over distributed data streams.

2. RELATED WORK

A well-studied problem is the monitoring of frequency counts over a single data stream [Alon et al. 1996; Arasu and Manku 2004; Charikar et al. 2002; Manku and Motwani 2002; Gibbons and Matias 1998]; however, these works do not address distributed environments.

Algorithms proposed in Dilman and Raz [2001] enable detecting when the sum of a distributed set of variables exceeds a predetermined threshold. However, the algorithms proposed in Dilman and Raz [2001] concentrate on monitoring the sum of a set of variables, whereas our algorithms enable monitoring arbitrary threshold functions over such variables. More recently, Keralapura et al. [2006] presented algorithms that adapt local thresholds when monitoring the sum of a set of variables, so that the communication cost is minimized. An interesting avenue for future work is making use of the techniques presented in Keralapura et al. [2006] when monitoring arbitrary threshold functions.

A scheme based on distributed hash tables, as well as a probabilistic coordinator based scheme for implementing distributed triggers, has been proposed in Jain et al. [2004]. This work analyzed a distributed trigger consisting of detecting when the sum of a distributed set of variables crosses a predetermined threshold. More recently, distributed triggers have been employed in order to detect anomalies in network traffic [Huang et al. 2006].

The algorithm proposed in Olston et al. [2003] enables a central coordinator to answer continuous queries designed to track the sum, average, or minimum of a distributed set of variables within a certain predetermined error margin. That work focused on minimizing the communication required for performing several concurrent monitoring tasks, whereas our work proposes algorithms for monitoring arbitrary threshold functions.

Babcock and Olston [2003] proposed an algorithm for finding the k largest aggregated values (for example the k largest frequency counters) over a set of distributed streams. Their algorithm employs a coordinator, which determines the initial set of the k largest aggregates, and sends each node a set of numerical constraints. Each node checks that the data received on its stream does not violate this constraint. As long as all the constraints are upheld, the list of top

k values is guaranteed to remain unchanged. In case a constraint is violated at one of the nodes, it notifies the coordinator, which performs a resolution process. The goal of the resolution process is to check if the list of the k largest values has changed, and to update the constraints at the nodes. The coordinator-based algorithm proposed in our work is similar to that algorithm in its use of numerical constraints, but its purpose is to monitor an arbitrary threshold function.

Algorithms proposed in Gibbons and Tirthapura [2001, 2002] enable estimating certain functions over a set of distributed streams, for example, the number of distinct elements in the streams, but this work does not address the monitoring of threshold functions. A preliminary version of this work appeared in Sharfman et al. [2006].

3. COMPUTATIONAL MODEL

Let $S = \{s_1, s_2, \dots, s_n\}$, be a set of n data streams, collected at nodes $P = \{p_1, p_2, \dots, p_n\}$. Let $\vec{v}_1(t), \vec{v}_2(t), \dots, \vec{v}_n(t)$ be d -dimensional real vectors derived from the streams (the value of these vectors varies over time). These vectors are called *local statistics vectors*. Let w_1, w_2, \dots, w_n be positive weights assigned to the streams.

The weight w_i assigned to the node p_i usually corresponds to the number of data items its local statistics vector is derived from. Assume, for example, that we would like to determine whether the frequency of occurrence of a certain data item in a set of streams is above a certain threshold value. In this case, the weight we assign to each node at time t is the number of data items received on the stream at time t (and $\vec{v}_i(t)$ is a scalar holding the frequency of occurrence of the item in the stream s_i). In this setup weights change over time. A variant of the problem stated above is for each node to maintain the frequency of occurrence of the item in the recent N_i data items received on the stream (this is known as working with a sliding window of size N_i). In that case, the weight assigned to each node is the size of its *sliding window*. In this setup weights do not change over time. For the sake of clarity, we assume at first that the weights are fixed, and that they are known to all nodes. Later, we modify our algorithms to handle weights that vary with time.

Let $\vec{v}(t) = \frac{\sum_{i=1}^n w_i \vec{v}_i(t)}{\sum_{i=1}^n w_i}$. $\vec{v}(t)$ is called the *global statistics vector*. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be an arbitrary function from the space of d -dimensional vectors to the reals. f is called the *monitored function*. We are interested in determining at any given time, t , whether or not $f(\vec{v}(t)) > r$, where r is a predetermined threshold value.

We present algorithms for two settings, a decentralized setting and a coordinator-based setting. Algorithms in both settings construct a vector called the *estimate vector*, denoted by $\vec{e}(t)$. The estimate vector is constructed from the local statistics vectors collected from the nodes at certain times, as dictated by the algorithms. The last statistics vector collected from the node p_i is denoted by \vec{v}'_i . Each node remembers the last statistics vector collected from it. The estimate vector is the weighted average of the latest statistics vectors collected from the nodes, that is, $\vec{e}(t) = \frac{\sum_{i=1}^n w_i \vec{v}'_i}{\sum_{i=1}^n w_i}$. From time to time, as dictated by the algorithm, an updated statistics vector is collected from one or more nodes, and

the estimate vector is updated. At any given time the estimate vector is known to all nodes.

In the decentralized setting, when the algorithm dictates that a statistics vector should be collected from a node, the node broadcasts the statistics vector to the rest of the nodes. Each node keeps track of the last statistics vector broadcast by every node, and locally calculates the estimate vector. In the coordinator-based setting, we designate a coordinator node and denote it by p_1 . The coordinator is responsible for collecting local statistics vectors from the nodes, calculating the estimate vector, and distributing it to the nodes.

In both settings, each node p_i maintains a parameter called the *statistics delta vector*. This vector is denoted by $\Delta\vec{v}_i(t)$. The statistics delta vector held by the node p_i is the difference between the current local statistics vector and the last statistics vector collected from the node, i.e., $\Delta\vec{v}_i(t) = \vec{v}_i(t) - \vec{v}'_i$.

In both settings, each node p_i also maintains a parameter called the *drift vector*. This vector is denoted by $\vec{u}_i(t)$. The drift vector is calculated differently in each setting. In the decentralized setting the drift vector is a displacement of $\Delta\vec{v}_i(t)$ in relation to the estimate vector,

$$\vec{u}_i(t) = \vec{e}(t) + \Delta\vec{v}_i(t). \quad (1)$$

The coordinator-based algorithm employs a mechanism for balancing the local statistics vectors of a subset of the nodes. Consider the case where at a certain time t the statistics delta vector in two equally weighted nodes, p_i and p_j , cancel each other out: that is, $\Delta\vec{v}_i(t) = -\Delta\vec{v}_j(t)$. We will see that balancing the local statistics vectors held by p_i and p_j can improve the efficiency of the algorithm. The coordinator facilitates this balancing by sending each node a *slack vector*, denoted by $\vec{\delta}_i$. The sum of the slack vectors sent to the nodes is $\vec{0}$. The drift vector held by each node is calculated as follows:

$$\vec{u}_i(t) = \vec{e}(t) + \Delta\vec{v}_i(t) + \frac{\vec{\delta}_i}{w_i}. \quad (2)$$

In the decentralized algorithm, nodes communicate by broadcasting messages. The cost of performing a broadcast varies according to the networking infrastructure at hand. In the worst case broadcasting a message to n nodes requires sending n point to point messages. While the decentralized algorithm remains highly efficient even in those settings, in practice, the cost of broadcasting a message is significantly lower. Some networking infrastructures, such as wireless networks and Ethernet based networks, support broadcasting at the cost of sending a single message. In other cases efficient broadcasting schemes have been developed that significantly reduce the cost of broadcasting.

We assume that communication links are reliable, that is, no messages are lost (otherwise standard methods for implementing reliability can be employed).

The notations presented in this section are summarized in Table I.

4. RUNNING EXAMPLE

Following is a detailed description of the spam filtering example presented in Section 1. We will use this example throughout the rest of the article to

Table I. Notation Summary

Notation	Description
$\bar{v}_i(t)$	Local statistics vector: the d -dimensional, time-varying data vector held by the i th node
w_i	Weight: a scalar weight assigned to the i th node
$\bar{v}(t)$	Global statistics vector: the weighted average of the local statistics vector
\bar{v}'_i	Sample vector: the value of the last local statistics vector collected from the i th node
$\bar{e}(t)$	Estimate vector: the weighted average of the sample vectors
$\Delta\bar{v}_i(t)$	Statistics delta vector: the difference between a node's local statistics vector and its sample vector
$\bar{\delta}_i$	Slack vector: an offset vector determined by the coordinator
$\bar{u}_i(t)$	Drift vector: the sum of the estimate vector and the statistics delta vector (includes the slack vector as well in the coordinator based algorithm, as specified by Equation (2))

demonstrate the concepts we present. Let p_1, p_2, \dots, p_n be n agents installed on n different mail servers. Let $M_i = \{m_{i,1}, m_{i,2}, \dots, m_{i,k}\}$ be the last k mail items received at the mail server installed on p_i , and let M denote the union of the last k mail items received at each one of the n mail servers, $M = \bigcup_{i=1}^n M_i$. Let X denote a set of mail items, let $Spam(X)$ be the set of mail items in X labeled as spam, and let $\overline{Spam}(X)$ be the set of mail items in X not labeled as spam. Let $Cont(X, f)$ be the set of mail items in X that contain the feature f , and let $\overline{Cont}(X, f)$ be the set of mail items in X that do not contain the feature f . Let the contingency table $C_{f,X}$ for the feature f over the set of mail items X be a 2×2 matrix, $C_{f,X} = \{c_{i,j}\}$, such that $c_{1,1} = \frac{|Cont(X,f) \cap Spam(X)|}{|X|}$, $c_{1,2} = \frac{|Cont(X,f) \cap \overline{Spam}(X)|}{|X|}$, $c_{2,1} = \frac{|\overline{Cont}(X,f) \cap Spam(X)|}{|X|}$, and $c_{2,2} = \frac{|\overline{Cont}(X,f) \cap \overline{Spam}(X)|}{|X|}$. C_{f,M_i} is called the *local contingency table* for the node p_i , and $C_{f,M}$ is called the *global contingency table*. Note that $C_{f,M} = \frac{\sum_{i=1}^n C_{f,M_i}}{n}$. We are interested in determining, for each feature f , whether the information gain over its global contingency table, denoted by $G(C_{f,M})$, is above or below a predetermined threshold r . The formula for information gain is given below:¹

$$G(C_{f,X}) = \sum_{i \in \{1,2\}} \sum_{j \in \{1,2\}} c_{i,j} \cdot \log \left(\frac{c_{i,j}}{(c_{i,1} + c_{i,2}) \cdot (c_{1,j} + c_{2,j})} \right).$$

According to the notations defined in Section 3, the local contingency table each node holds is the local statistics vector, and the global contingency table is the global statistics vector. Note that the answer to the threshold function query *cannot* be derived from the value of the monitored function on data from each individual stream. Consider, for example, a spam filtering system consisting of two streams, with a threshold value of 0.5. The first node may hold a contingency table $C_{f,M_1} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$, resulting in $G(C_{f,M_1}) = 0$, and the second node may hold a contingency table $C_{f,M_2} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$, resulting in $G(C_{f,M_2}) = 0$. As we can see, the gain calculated on each individual stream is 0, and thus below the threshold value, but the gain on the global contingency table for

¹If $c_{i,j} = 0$ then $c_{i,j} \cdot \log \left(\frac{c_{i,j}}{(c_{i,1} + c_{i,2}) \cdot (c_{1,j} + c_{2,j})} \right)$ is defined as 0.

f , $C_{f, M_1 \cup M_2} = \frac{C_{f, M_1} + C_{f, M_2}}{2} = \begin{pmatrix} 0.5 & 0 \\ 0 & 0.5 \end{pmatrix}$, is $G(C_{f, M_1 \cup M_2}) = 1$, and thus above the threshold value. Note that this behavior does not occur when monitoring frequencies of occurrence of items over distributed streams, that is, if the frequency of occurrence of a certain item in all the streams is below a predetermined threshold, then the frequency of occurrence of that item over the union of the streams is below the threshold as well.

5. GEOMETRIC INTERPRETATION

At the heart of our approach is the ability to decompose the monitoring task into local constraints on streams. As data arrives on the streams, each node verifies that the local constraint on its stream has not been violated. We will show that as long as none of these constraints has been violated, the query result is guaranteed to remain unchanged, and thus no communication is required. As demonstrated in Section 1, this cannot be done solely by observing the value of the monitored function on each stream. Therefore, we have defined in Section 3 an estimated global statistics vector, called the *estimate vector*, which is known to all nodes. The estimate vector is said to be correct at a given time if the value of the monitored function on the estimate vector and the value of the monitored function on the global statistics vector at that time (this value is unknown to any single node) are on the same side of the threshold. Given an initially correct estimate vector, our goal is to set local constraints on each stream such that as long as no constraints have been violated, the estimate vector remains correct, and thus no communication is required. The method for decomposing the monitoring task is based on the following, easily verifiable observation—at any given time the weighted average of the drift vectors held by the nodes is equal to the global statistics vector,

$$\frac{\sum_{i=1}^n w_i \vec{u}_i(t)}{\sum_{i=1}^n w_i} = \vec{v}(t). \quad (3)$$

We refer to Property (3) as the *convexity property* of the drift vectors. The geometric interpretation of Property (3) is that the global statistics vector is in the convex hull of the drift vectors held by the nodes,

$$\vec{v}(t) \in \text{Conv}(\vec{u}_1(t), \vec{u}_2(t), \dots, \vec{u}_n(t)). \quad (4)$$

This observation enables us to take advantage of Theorem 1 in order to decompose the monitoring task.

THEOREM 1. *Let $\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n \in \mathbb{R}^d$ be a set of vectors in \mathbb{R}^d . Let $\text{Conv}(\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n)$ be the convex hull of $\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n$. Let $B(\vec{x}, \vec{y}_i)$ be a ball centered at $\frac{\vec{x} + \vec{y}_i}{2}$ and with a radius of $\|\frac{\vec{x} - \vec{y}_i}{2}\|_2$ that is, $B(\vec{x}, \vec{y}_i) = \{\vec{z} \mid \|\vec{z} - \frac{\vec{x} + \vec{y}_i}{2}\|_2 \leq \|\frac{\vec{x} - \vec{y}_i}{2}\|_2\}$; then $\text{Conv}(\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n) \subset \bigcup_{i=1}^n B(\vec{x}, \vec{y}_i)$.*

Theorem 1 is used to bound the convex hull of $n + 1$ vectors in \mathbb{R}^d by the union of n d -dimensional balls. In our case it is used to bound the convex hull of the

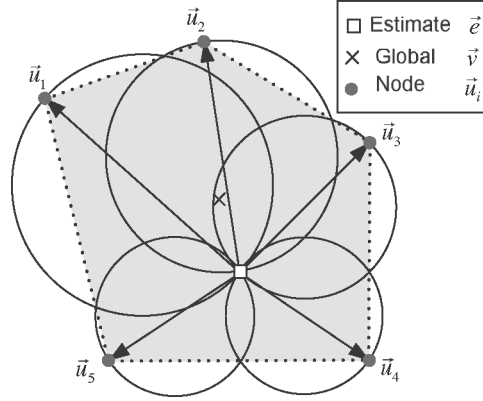


Fig. 1. Illustration of Theorem 1. The drift vectors held by five nodes and the balls constructed by them are depicted. The convex hull of the drift vectors is highlighted in gray. As stated by the theorem, the union of the balls bounds the convex hull.

estimate vector and the drift vectors, that is, $\text{Conv}(\vec{e}(t), \vec{u}_1(t), \vec{u}_2(t), \dots, \vec{u}_n(t))$, by a set of n balls, where each ball is constructed independently by one of the nodes. Each node, p_i , constructs a ball $B(\vec{e}(t), \vec{u}_i(t))$, which is centered at $\frac{\vec{e}(t) + \vec{u}_i(t)}{2}$, and has a radius of $\|\frac{\vec{e}(t) - \vec{u}_i(t)}{2}\|$. Note that at any given time each node has all the information required to independently construct its ball. Theorem 1 states that $\text{Conv}(\vec{e}(t), \vec{u}_1(t), \vec{u}_2(t), \dots, \vec{u}_n(t)) \subset \bigcup_i B(\vec{e}(t), \vec{u}_i(t))$.

The application of Theorem 1 is illustrated in Figure 1, which depicts a setup comprised of five nodes, each holding a statistics vector, $\vec{v}_i(t) \in \mathbb{R}^2$. The drift vectors held by the nodes ($\vec{u}_1(t), \dots, \vec{u}_5(t)$), the global statistics vector $\vec{v}(t)$, and the estimate vector $\vec{e}(t)$ are depicted, as are the balls constructed by the nodes. The convex hull of the drift vectors is highlighted in gray, and one can see that, as the theorem states, the area defined by the convex hull is bounded by the set of balls.

5.1 Local Constraints

The local constraint on each stream is set as follows: the monitored function f and threshold r can be seen as inducing a coloring over \mathbb{R}^d . The vectors $\{\vec{x} \mid f(\vec{x}) > r\}$ are said to be *green*, while the vectors $\{\vec{y} \mid f(\vec{y}) \leq r\}$ are said to be *red*. The local constraint each node maintains is to check whether the ball $B(\vec{e}(t), \vec{u}_i(t))$ (the ball centered at $\frac{\vec{e}(t) + \vec{u}_i(t)}{2}$, and having a radius of $\|\frac{\vec{e}(t) - \vec{u}_i(t)}{2}\|$) is monochromatic, that is, whether all the vectors contained in the ball have the same color. Testing for monochromaticity is done by finding the maximal and minimal values of f in the ball. This is done locally at each node hence has no effect on the communication load.

If all the local constraints are upheld, the estimate vector is correct: because all the balls contain the estimate vector, and all the balls are monochromatic, the set of vectors defined by the union of all the balls is monochromatic as well. Since the union of all the balls contains the convex hull of the drift vectors and the estimate vector ($\text{Conv}(\vec{e}(t), \vec{u}_1(t), \vec{u}_2(t), \dots, \vec{u}_n(t))$), and according

to Equation (4) the global statistics vector is contained in the convex hull of the drift vectors, the estimate vector and the global statistics vector have the same color. Therefore, they are on the same side of the threshold, that is, the estimate vector is correct.

5.2 Proof of Theorem 1

Following is the proof of Theorem 1. The proof uses the following variant of Caratheodory's theorem [Berkovitz 2002]:

THEOREM 2. *Let $\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n \in \mathbb{R}^d$ be a set of vectors in \mathbb{R}^d . Let $\text{Conv}(\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n)$ be the convex hull of $\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n$. Any vector $\vec{z} \in \text{Conv}(\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n)$ can be expressed as a convex combination of \vec{x} , and at most d members of $\{\vec{y}_1, \vec{y}_2, \dots, \vec{y}_n\}$.*

The proof will use the following Lemma:

LEMMA 1. *Let $\vec{x}_1, \vec{x}_2 \dots \vec{x}_d$ be d vectors in \mathbb{R}^d . A rotation matrix R exists such that the last coordinate of $R\vec{x}_i$ is identical for all $1 \leq i \leq d$.*

PROOF. If \vec{r}^T denotes the last row of R , we must have $\vec{r}^T \vec{x}_1 = \vec{r}^T \vec{x}_2 = \dots = \vec{r}^T \vec{x}_d$ for all $1 \leq i \leq d$. Let us differentiate between two cases. If \vec{x}_i are linearly dependent, their span lies in a subspace of dimension less than d ; hence there's a nonzero vector which is orthogonal to all of them; take \vec{r}^T to be equal to that vector. If they are independent, look at the system of d linear equations $\vec{r}^T \vec{x}_1 = 1, \dots, \vec{r}^T \vec{x}_d = 1$. Its coefficient matrix is composed of the \vec{x}_i ; hence it has full rank, and thus admits a solution, which we'll take to be \vec{r}^T . In both cases, normalize \vec{r}^T so that it has unit length; the requirement $\vec{r}^T \vec{x}_1 = \vec{r}^T \vec{x}_2 = \dots = \vec{r}^T \vec{x}_d$ obviously still holds.

Last, we need to construct a rotation matrix R having \vec{r}^T as its last row. Since \vec{r}^T is a unit vector, we can use the Gram-Schmidt process to complete it to an orthonormal basis of \mathbb{R}^d by adding to it $d - 1$ vectors of unit length such that all d vectors are orthogonal to each other. Take R to consist of this basis (here we use the fact that R is a rotation matrix iff its rows form an orthonormal basis). \square

PROOF OF THEOREM 1. The theorem is proven by induction over d . The base of the induction is $d = 1$. Proving the base of the induction is trivial ($\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n$ are real numbers and both $\text{Conv}(\vec{x}, \vec{y}_i)$ and $B(\vec{x}, \vec{y}_i)$ are the segment $[x, y_i]$).

The induction step is proven as follows: according to Theorem 2, since any vector $\vec{z} \in \text{Conv}(\vec{x}, \vec{y}_1, \vec{y}_2, \dots, \vec{y}_n)$ can be expressed as a convex combination of \vec{x} and a set of at most d members of $\{\vec{y}_1, \vec{y}_2, \dots, \vec{y}_n\}$, it is sufficient to show that for an arbitrary set of vectors, $\{\vec{x}, \vec{v}_1, \vec{v}_2, \dots, \vec{v}_d\} \in \mathbb{R}^d$,

$$\text{Conv}(\vec{x}, \vec{v}_1, \vec{v}_2, \dots, \vec{v}_d) \subset \bigcup_{i=1}^d B(\vec{x}, \vec{v}_i). \quad (5)$$

Furthermore, we observe that any vector $\vec{z} \in \text{Conv}(\vec{x}, \vec{v}_1, \vec{v}_2, \dots, \vec{v}_d)$ is a convex combination of some vector $\vec{z}' \in \text{Conv}(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_d)$ and \vec{x} . Therefore, it is

sufficient to show that

$$\text{Conv}(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_d) \subset \bigcup_{i=1}^d B(\vec{x}, \vec{v}_i). \quad (6)$$

Equation (6) implies Equation (5) since \vec{z} is a convex combination of $\vec{z}' \in \text{Conv}(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_d)$ and \vec{x} , and since, according to Equation (6), $\text{Conv}(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_d) \subset \bigcup_{i=1}^d B(\vec{x}, \vec{v}_i)$, \vec{z}' belongs to at least one of the balls $B(\vec{x}, \vec{v}_1), B(\vec{x}, \vec{v}_2), \dots, B(\vec{x}, \vec{v}_d)$, say $B(\vec{x}, \vec{v}_j)$. Since by definition \vec{x} also belongs to $B(\vec{x}, \vec{v}_j)$, \vec{z} also belongs to $B(\vec{x}, \vec{v}_j)$. This is because \vec{z} is a convex combination of \vec{z}' and \vec{x} , which both belong to $B(\vec{x}, \vec{v}_j)$, which is a convex set.

Since for any rotation matrix R and vectors $\{q_1 \dots q_n\}$, $\text{Conv}\{R(q_1) \dots R(q_n)\} = R(\text{Conv}\{q_1 \dots q_n\})$, and Lemma 1 guarantees that a rotation matrix R can be found such that the last coordinate of $R\vec{v}_1, R\vec{v}_2, \dots, R\vec{v}_d$ is identical, we assume without loss of generality that the vectors $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_d$ lie on a $d - 1$ dimensional hyperplane, P , which consists of vectors whose last coordinate is identical. It is easy to show that the intersection of $B(\vec{x}, \vec{v}_1), B(\vec{x}, \vec{v}_2), \dots, B(\vec{x}, \vec{v}_d)$ with P yields the set of $d - 1$ dimensional balls, $B(\vec{x}_p, \vec{v}_1), B(\vec{x}_p, \vec{v}_2), \dots, B(\vec{x}_p, \vec{v}_d)$, where \vec{x}_p is the projection of \vec{x} on the plane P . Note that, for every i , the $d - 1$ dimensional ball $B(\vec{x}_p, \vec{v}_i)$ is contained in the d dimensional ball $B(\vec{x}, \vec{v}_i)$ ($B(\vec{x}_p, \vec{v}_i) \subset B(\vec{x}, \vec{v}_i)$). Since the vectors $\vec{x}_p, \vec{v}_1, \vec{v}_2, \dots, \vec{v}_d$ lie on the $d - 1$ dimensional plane P , then, according to the induction hypothesis, $\text{Conv}(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_d) \subset \bigcup_{i=1}^d B(\vec{x}_p, \vec{v}_i)$. Since for every i , $B(\vec{x}_p, \vec{v}_i)$ is contained in $B(\vec{x}, \vec{v}_i)$, then $\text{Conv}(\vec{v}_1, \vec{v}_2, \dots, \vec{v}_d) \subset \bigcup_{i=1}^d B(\vec{x}, \vec{v}_i)$. We have proved Equation (6), and thus concluded the proof. \square

6. DISTRIBUTED MONITORING

In this section we present algorithms that are based on the geometric method for decomposing the monitoring task into local constraints on the streams. After presenting the algorithms, we describe how they can be tuned to relax performance requirements in favor of reducing communication load, and how they can be modified to support time-varying weights.

6.1 The Decentralized Algorithm

Following is a simple, broadcast-based algorithm for monitoring threshold functions: each node maintains a copy of the last statistics vector sent by each of the nodes. The initialization stage consists of every node broadcasting its initial statistics vector. Upon receipt of all the initialization messages, each node calculates the estimate vector, $\vec{e}(t)$. Then, as more data arrives on the stream, each node can check its local constraint according to the estimate vector and its drift vector. If a local constraint is violated at a node, p_i , the node broadcasts a message of the form $\langle i, \vec{v}_i(t) \rangle$, containing its identifier and its local statistics vector at the time. The broadcasting node updates its \vec{v}'_i parameter and recalculates the estimate vector. Upon receiving a broadcast message from a node, p_i , each node updates its \vec{v}'_i parameter and recalculates the estimate vector.

If all the local constraints are upheld, Theorem 1 guarantees the correctness of the estimate vector (enabling every node to locally calculate the value of

Algorithm 1 The decentralized algorithm

Initialization: at a node p_i

- Broadcast a message containing the initial statistics vector and update \vec{v}'_i to hold the initial statistics vector. Upon receipt of messages from all the nodes, calculate the estimate vector $(\vec{e}(t))$.

Processing stage at a node p_i :

- Upon arrival of new data on the local stream, recalculate $\vec{v}_i(t)$, and $\vec{u}_i(t)$, and check if $B(\vec{e}(t), \vec{u}_i(t))$ remains monochromatic. If not, broadcast the message $\langle i, \vec{v}_i(t) \rangle$ and update \vec{v}'_i to hold $\vec{v}_i(t)$.
 - Upon receipt of a new message $\langle j, \vec{v}_j(t) \rangle$, update \vec{v}'_j to hold $\vec{v}_j(t)$, recalculate $\vec{e}(t)$, and check if $B(\vec{e}(t), \vec{u}_i(t))$ is monochromatic. If $B(\vec{e}(t), \vec{u}_i(t))$ is not monochromatic, broadcast the message $\langle i, \vec{v}_i(t) \rangle$ and update \vec{v}'_i to hold $\vec{v}_i(t)$.
-

the threshold function). After a node broadcasts a message, its local constraint is upheld (because the ball it constructs has a radius of 0, and therefore is monochromatic). If a local constraint has been violated, at worst all n nodes (but possibly fewer) will broadcast a message before all the local constraints are upheld again.

A formal description can be found in Algorithm 1.

6.2 The Coordinator-Based Algorithm

Local constraints are also used in the coordinator-based algorithm, but the coordinator is responsible for calculating the estimate vector, maintaining its correctness, and distributing it to the other nodes. In the decentralized algorithm the violation of a constraint on one node requires communicating with all the rest of the nodes (a broadcast message is sent). While this may be a good solution in setups where the nodes are closely coupled, in other cases we can further reduce the communication load by introducing a coordinator. The presence of a coordinator enables us to resolve a violation at a node by communicating with only a subset of the nodes, as opposed to communicating with all the nodes as required in the decentralized algorithm. Consider, for example, a set of equally weighted nodes monitoring the function $f(x) = (x - 5)^2$ (a function over a single dimensional statistics vector), and a threshold value of $r = 9$. Say that at time t the estimate vector is $\vec{e}(t) = 5$. Note that, since $f(\vec{e}(t)) = 0 < r$, any drift vector in the range $[2, 8]$ satisfies the local constraint at the node. Let us assume that the drift vector at the coordinator, p_1 , is $\vec{u}_1(t) = 4$, and the constraints at all n nodes are satisfied except for p_2 , which holds the drift vector $\vec{u}_2(t) = 1$. In the decentralized algorithm, since the constraint at p_2 has been violated, it would have broadcast its statistics vector to all n nodes. However, the constraint violation at p_2 can be resolved by setting the drift vector at both p_1 and p_2 to the average of the drift vectors on both nodes, that is, by setting $\vec{u}_1(t) = \vec{u}_2(t) = \frac{\vec{u}_1(t) + \vec{u}_2(t)}{2} = 2.5$. After this averaging operation, drift vectors on both p_1 and p_2 are within the range $[2, 8]$, and thus all the local constraints are upheld. Note that this action preserves the convexity property of the drift

vector (Property 3). The act of averaging out a subset of drift vectors in order to resolve a violated constraint is called a balancing process.

In order to facilitate the balancing of vectors, every node p_i holds a slack vector denoted by $\vec{\delta}_i$, as defined in Section 3. The slack vector is first normalized by dividing it by the weight assigned to the node. Then it is added to the drift vector as specified in Equation (2). The coordinator is responsible for ensuring that the sum of all slack vectors is $\vec{0}$, thus maintaining the convexity property of the drift vectors (Property 3).

To initialize the algorithm, each node sends its initial statistics vector to the coordinator. Initially the slack vector held by each node is set to $\vec{0}$. The coordinator calculates the estimate vector and sends it to the rest of the nodes. As more data arrives on a node's stream, the node checks its local constraint. If a local constraint is violated at one of the nodes, it notifies the coordinator by sending it a message containing its current drift vector and its current statistics vector. The coordinator first tries to resolve the constraint violation by executing a balancing process.

During the balancing process the coordinator tries to establish a group of nodes (called the *balancing group* and denoted by P'), such that the average of the drift vectors held by the nodes in the balancing group (called the balanced vector and denoted by \vec{b}), creates a monochromatic ball with the estimate vector, that is, such that $B(\vec{e}(t), \vec{b})$ is monochromatic. The balanced vector is calculated as follows:

$$\vec{b} = \frac{\sum_{p_i \in P'} w_i \vec{u}_i(t)}{\sum_{p_i \in P'} w_i} \quad (7)$$

The balancing process proceeds as follows: when a node p_i notifies the coordinator that its local constraint has been violated, it appends its drift vector and its current statistics vector to the message. The coordinator constructs a balancing group consisting of p_i and itself. It then checks if the ball defined by the balanced vector, $B(\vec{e}(t), \vec{b})$, is monochromatic. If $B(\vec{e}(t), \vec{b})$ is not monochromatic, the coordinator randomly selects a node that is not in the balancing group, and requests it to send its drift vector and local statistics vector. Then it adds that new node to the balancing group and rechecks $B(\vec{e}(t), \vec{b})$. The process is performed iteratively until either $B(\vec{e}(t), \vec{b})$ is monochromatic or the balancing group contains all the nodes. If the coordinator established a balancing group such that $B(\vec{e}(t), \vec{b})$ is monochromatic, the balancing process is said to have succeeded. In this case the coordinator sends each node in the balancing group an adjustment to its slack vector. This causes the drift vectors held by all nodes in the balancing group to be equal to \vec{b} . The adjustment to the slack vector sent to each node $p_i \in P'$ is denoted by $\Delta \vec{\delta}_i$, and is calculated as follows:

$$\Delta \vec{\delta}_i = w_i \vec{b} - w_i \vec{u}_i(t).$$

After receiving the slack vector adjustment, each node simply adds the adjustment to the current value, that is, $\vec{\delta}_i \leftarrow \vec{\delta}_i + \Delta \vec{\delta}_i$. One can easily verify that after a successful balancing process the sum of all slack vectors remains $\vec{0}$, and

the drift vector held by each node in the balancing group is \vec{b} , thus resolving the original constraint violation.

If the balancing process has failed (i.e., the balancing group contains all the nodes, and $B(\vec{e}(t), \vec{b})$ is not monochromatic), the coordinator calculates a new estimate vector (according to the updated statistics vectors sent by the nodes) and sends it to all the nodes. Upon receipt of the new estimate vector, the nodes set their slack vectors to $\vec{0}$ and modify their \vec{v}'_i parameter to hold the value of the statistics vector they sent to the coordinator during the balancing process, thus resolving the original constraint violation.

In order to implement the algorithm, the following messages must be defined:

- <INIT, \vec{v}_i > Used by nodes to report their initial statistics vector to the coordinator in the initialization stage.
- <REQ> Used by the coordinator during the balancing process to request that a node send its statistics vector and drift vector.
- <REP, \vec{v}_i, \vec{u}_i > Used by nodes to report information to the coordinator when a local constraint has been violated, or when the coordinator requests information from the node.
- <ADJ-SLK, $\Delta\vec{\delta}_i$ > Used by the coordinator to report slack vector adjustments to nodes after a successful balancing process.
- <NEW-EST, \vec{e} > Used by the coordinator to report to the nodes a new estimate vector.

A formal description is given in Algorithm 2.

6.3 Relaxing the Precision Requirements

A desired tradeoff when monitoring threshold functions is between accuracy and communication load. In some cases an approximate value of the threshold function is sufficient, that is, the correct value of the threshold function is required only if the value of the monitored function is significantly far from the threshold. In other words, if ε is a predetermined error margin, and if $f(\vec{v}(t)) > r + \varepsilon$ or $f(\vec{v}(t)) \leq r - \varepsilon$, we require that the estimate vector, $\vec{e}(t)$, be correct, but we do not require it if $r - \varepsilon < f(\vec{v}(t)) \leq r + \varepsilon$.

Consider the feature monitoring example given in Section 1. Say we would like to select all the features whose information gain is above 0.05. Obviously, it is important to select a feature whose information gain is significantly high, and not to select a feature whose information gain is significantly low. For example, it is important to select a feature whose information gain score is 0.1, and not to select a feature whose information gain score is 0.01. Including or excluding features whose information gain score is very close to the threshold value; for example, a feature whose information gain score is 0.048, will probably not have a significant effect on the quality of the selected feature set, while the cost of monitoring such features is expected to be high, since their information gain is expected to fluctuate around the threshold value. Therefore we can significantly improve the efficiency of our monitoring algorithms if we set some error margin, say 0.005. In other words, features that are currently selected will be

Algorithm 2 The coordinator-based algorithm

Initialization:

—Send an INIT message to the coordinator, set \vec{v}' to hold the initial statistics vector, and set the slack vector to $\vec{0}$. Upon receipt of messages from all nodes, the coordinator calculates the estimate vector and informs the nodes via a NEW-EST message.

Processing stage at an ordinary node p_i :

—Upon arrival of new data on a node's local stream, recalculate $\vec{v}_i(t)$ and $\vec{u}_i(t)$, and check if $B(\vec{e}(t), \vec{u}_i(t))$ remains monochromatic. If not, send a $\langle \text{REP}, \vec{v}_i(t), \vec{u}_i(t) \rangle$ message to the coordinator, and wait for either a NEW-EST or an ADJ-SLK message.

—Upon receipt of a REQ message, send a $\langle \text{REP}, \vec{v}_i(t), \vec{u}_i(t) \rangle$ message to the coordinator and wait for either a NEW-EST or ADJ-SLK message.

—Upon receipt of a NEW-EST message, update the estimate vector ($\vec{e}(t)$) to the value specified in the message, set the value of \vec{v}' to the statistics vector sent to the coordinator, and set the slack vector to $\vec{0}$.

—Upon receipt of an ADJ-SLK message, add the value specified in the message to the value of the slack vector ($\vec{\delta}_i \leftarrow \vec{\delta}_i + \Delta \vec{\delta}_i$).

Processing stage at the coordinator:

—Upon arrival of new data on the local stream, recalculate $\vec{v}_1(t)$ and $\vec{u}_1(t)$, and check if $B(\vec{e}(t), \vec{u}_1(t))$ remains monochromatic. If not, initiate a balancing process, setting the balancing group to $P' = \{ \langle 1, \vec{v}_1(t), \vec{u}_1(t) \rangle \}$.

—Upon receipt of a REP message from the node p_i , initiate a balancing process, setting the balancing group to $P' = \{ \langle 1, \vec{v}_1(t), \vec{u}_1(t) \rangle, \langle i, \vec{v}_i, \vec{u}_i \rangle \}$.

Balancing process at the coordinator:

- (1) Calculate balanced vector, \vec{b} , according to Equation (7). If the ball $B(\vec{e}(t), \vec{b})$ is monochromatic goto (2), otherwise goto (3).
- (2) For each item in the balancing group, $\langle i, \vec{v}_i, \vec{u}_i \rangle$, calculate the slack vector adjustment, $\Delta \vec{\delta}_i = w_i \vec{b} - w_i \vec{u}_i(t)$, send p_i a $\langle \text{ADJ-SLK}, \Delta \vec{\delta}_i \rangle$ message, and then exit the Balancing Process.
- (3) If there are nodes not contained in the balancing group, select one of these nodes at random, and send it a REQ message. Upon receipt of the REP message, add the node to the balancing group and goto (1). Otherwise calculate a new estimate vector (based on the \vec{v}_i values received from all the nodes), send a NEW-EST message to all nodes, and exit the balancing process.

removed from the set of selected features only when their information gain falls below 0.045, and features that are currently not selected will be added to the set of selected feature only if their information gain rises above 0.055.

Our algorithm can be easily tuned to relax the precision requirements by an error margin of ε as follows: instead of working with a single coloring, induced by the monitored function f and the threshold value r , two sets of coloring are defined, one induced by the monitored function f and the threshold value $r + \varepsilon$, and a second induced by the monitored function f and the threshold value $r - \varepsilon$. Whenever the original algorithm checks whether a ball is monochromatic, then, if $f(\vec{v}(t)) \leq r$, the modified algorithm will check whether the ball is monochromatic according to the first coloring (the one induced by f and $r + \varepsilon$).

If $f(\vec{v}(t)) > r$, the modified algorithm will check whether the ball is monochromatic according to the second coloring (the one induced by f and $r - \varepsilon$). This ensures that, if all the balls are in the range defined by $\{\vec{x} | r - \varepsilon < f(\vec{x}) \leq r + \varepsilon\}$, no messages are transmitted.

6.4 Handling Time-Varying Weights

Up to this point we have assumed that the weights assigned to nodes are fixed, such as when the weights are the size of sliding windows used for collecting data from the streams. We now address cases where weights assigned to nodes may vary with time, as when a node's weight at a given time is the number of data items received on its stream so far.

We next describe the required modifications to the algorithms in order to ensure their correctness in a setup where weights vary with time. In such a setup we denote the weight assigned to the node p_i at time t by $w_i(t)$. Each message in the original algorithms is modified by appending $w_i(t)$ to it. Along with the last vector broadcast by each of the other nodes (\vec{v}'_i), the nodes in the decentralized algorithm keep track of the last broadcast weight, denoted by w'_i . Nodes calculate $\vec{e}_i(t)$, $\Delta\vec{v}_i(t)$, and $\vec{u}_i(t)$ as follows:

$$\vec{e}_i(t) = \frac{\sum_{i=1}^n w'_i \vec{v}'_i}{\sum_{i=1}^n w'_i},$$

$$\Delta\vec{v}_i(t) = \frac{w_i(t)\vec{v}_i(t) - w'_i\vec{v}'_i(t) - (w_i(t) - w'_i)\vec{e}_i(t)}{w_i(t)}.$$

In the decentralized algorithm, the drift vector is calculated by

$$\vec{u}_i(t) = \vec{e}_i(t) + \Delta\vec{v}_i(t)$$

and in the coordinator-based algorithm, by

$$\vec{u}_i(t) = \vec{e}(t) + \Delta\vec{v}_i(t) + \frac{\vec{\delta}_i}{w_i(t)}.$$

In the coordinator-based algorithm, the balanced vector and the slack vector adjustments are calculated according to the weights appended to the messages:

$$\vec{b} = \frac{\sum_{i=1}^k w_i(t)\vec{u}_i(t)}{\sum_{i=1}^k w_i(t)},$$

$$\Delta\vec{\delta}_i = w_i(t)\vec{b} - w_i(t)\vec{u}_i(t).$$

Note that, if the weights are fixed, $\vec{e}_i(t)$, $\Delta\vec{v}_i(t)$, and $\vec{u}_i(t)$ hold the same values they hold in the original algorithms. Furthermore, one can easily verify that the new definitions of these parameters maintain Equation (3), that is,

$\vec{v}(t) = \frac{\sum_{i=1}^n w_i(t)\vec{u}_i(t)}{\sum_{i=1}^n w_i(t)}$, and thus maintain the correctness of the algorithm.

7. DETERMINING MONOCHROMATICITY

In this section, we address the computational problem of determining whether a sphere is monochromatic. This is equivalent to testing whether the sphere

intersects the threshold surface $\{p \mid f(p) = t\}$, where f is the monitored function and t the threshold.

The most direct solution is to compute the distance from the sphere's center to the threshold surface. This is a constrained optimization problem which has been extensively studied, for example, in the context of implicit surface fitting. For example, in Sullivan et al. [1994] the running time (modified to current-day computers) for computing the distance of a point in three-dimensional Euclidean space from an implicit surface defined by a fourth-degree polynomial is, on the average, about 10^{-3} s. Specifically for polynomials, recent algorithms were presented for performing global optimizations over polynomials with polynomial constraints [Parrilo 2003; Lasserre 2000]. These algorithms can be used in order to directly calculate the distance of a given vector from a threshold surface when the monitored function is a polynomial.

If one allows for some rather modest preprocessing, faster algorithms can be designed. As opposed to the problem of surface fitting, we have the advantage that the monitored functions are fixed. Therefore, the threshold surface can be computed in advance. Moreover, the approximate distance of every point from the threshold surface (the *distance transform*) can be estimated very accurately by computing it on a dense grid, using algorithms that are linear in the size of the grid. This holds also for nonpolynomial functions, such as the information gain.

To get a feeling for the time required for this preprocessing—which has only to be performed once—we refer to the popular recursive subdivision approach; for example, in Varadhan et al. [2004] running times of about 1 min on the average are reported for computing the threshold surface of very complicated functions. The distance transform can also be computed quickly; for example, in Siddiqi et al. [2002] a running time of 26 s is reported for computing the distance transform for a $128 \times 128 \times 128$ array.

We reiterate that this preprocessing has to be performed only once; after it's done, each node keeps a copy of the distance array, and then testing for monochromaticity is extremely fast.

8. PERFORMANCE ANALYSIS

We would like to determine how the various parameters of the monitoring problem affect the communication load generated by the proposed algorithms. In order to do so we present a simplified model of our algorithm and analyze the probability that a constraint violation will occur at a node. Since a constraint violation is the trigger for communications in both algorithms, this analysis should provide indications regarding the generated communication load.

Generally speaking, the dominant factor affecting the performance of the algorithms is the average distance of the estimate vector from the set of vectors for which the value of the monitoring function equals the threshold value. More formally, let the threshold set defined by the monitoring function f and threshold value r be the set of vectors for which the value of the threshold function equals the threshold value. Let the threshold set be denoted by $T(f, r)$, that is,

$$T(f, r) = \{\vec{x} \mid f(\vec{x}) = r\}.$$




Fig. 2. The colorings induced by two sets of monitored functions and threshold values. (a) The coloring induced by $f_1 = \sin(2\sqrt{x^2 + y^2}) \geq 0$, and (b) the coloring induced by $f_2 = \frac{1}{1+e^{-x}} + \frac{1}{1+e^{-y}} \geq 0.75$.

Let the distance of a vector \vec{x} from the threshold set $T(f, r)$, denoted as $dist(\vec{x}, f, r)$, be the minimum distance of \vec{x} from any point in $T(f, r)$ that is,

$$dist(\vec{x}, f, r) = \min(\|\vec{y} - \vec{x}\| \mid f(\vec{y}) = r).$$

The farther the estimate vector is, at a given time, from the threshold set, the more the local statistics vectors can change without violating local constraints. Therefore, a greater average distance of the estimate vector from the threshold set will result in a greater reduction in communications.

The average distance of the estimate vector from the threshold set is affected by many parameters. To begin with, it is affected by the coloring induced by the monitored function and the threshold value. Figure 2 illustrates the coloring induced by two sets of a monitored function and a threshold value. Figure 2(a) illustrates the coloring induced by the function $f_1(x, y) = \sin(2\sqrt{x^2 + y^2})$ and the threshold value 0, and Figure 2(b) illustrates the coloring induced by the function $f_2(x, y) = \frac{1}{1+e^{-x}} + \frac{1}{1+e^{-y}}$ and the threshold value 0.75 ($f_2(x, y)$ is a simple two-layer neural net).

It is clear that the distance of any point in \mathbb{R}^2 from the threshold set defined by f_1 and 0 cannot be greater than $\frac{\pi}{4} \approx 0.785$. Therefore, the average distance of the estimate vector from the threshold set in this case is bounded from above by 0.785, thus yielding a relatively low reduction in communications when monitored by our algorithms. However, the maximum distance of a point in \mathbb{R}^2 from the threshold set defined by f_2 and 0.75 is unbounded, and the dominating factor affecting the performance of our algorithms in this case is the nature of the data received on the streams.

In order to analyze how our algorithms are affected by the nature of the data on the streams, we consider periods during which this data is stationary (this fact, however, is not known to any of the nodes). More formally, we assume that each stream item is a d -dimensional vector, where the j th component is independently drawn from a random variable denoted by X_j with a defined expectancy and variance, denoted by $E[X_j]$ and $V[X_j]$, respectively. We assume the system consists of n nodes, and that each node holds a




Fig. 3. The coloring induced by the function $f_2 = \frac{1}{1+e^{-x}} + \frac{1}{1+e^{-y}}$ and the threshold value 0.75, together with the expected global statistics vector, its distance from the threshold set, and a local statistics vector, that is contained within the distance sphere. One can see that the ball $B(E[\bar{v}(t)], \bar{v}_i(t))$ is fully contained in the distance sphere.

sliding window of N items. We denote the last N items received on the stream monitored by p_i as $\bar{v}_{i,1}, \bar{v}_{i,2}, \dots, \bar{v}_{i,N}$, and the components of a vector as follows: $\bar{v}_{i,k} = (v_{i,k}^{(1)}, v_{i,k}^{(2)}, \dots, v_{i,k}^{(d)})$. The local statistics vector held by a node is the average of the items contained in its sliding window, and the global statistics vector is the average of the items contained in the sliding windows held by all the nodes, that is,

$$\bar{v}_i(t) = \frac{\sum_{k=1}^N \bar{v}_{i,k}}{N} = \sum_{k=1}^N \frac{\bar{v}_{i,k}}{N}, \quad (8)$$

$$\bar{v}(t) = \frac{\sum_{i=1}^n \sum_{k=1}^N \bar{v}_{i,k}}{N \cdot n}. \quad (9)$$

It is easy to see that the expected value for the global statistics vector and each local statistics vector is $E[\bar{v}(t)] = E[\bar{v}_i(t)] = (E[X_1], E[X_2], \dots, E[X_d])$.

Figure 3 depicts the coloring induced by f_2 and the threshold value 0.75, the expected global statistics vector, and the distance of the expected global statistics vector from the threshold set. The expected global statistics vector and its distance from the threshold set define a sphere called the *distance sphere*. We denote the distance of the expected global statistics vector from the threshold set by D_{global} , that is, $D_{global} = \text{dist}(E[\bar{v}(t)], f, r)$.

We present the following simplified model of our algorithms: we assume that the estimate vector holds the value of the expected global statistics vector, that is, $\bar{e}(t) = E[\bar{v}(t)] = (E[X_1], E[X_2], \dots, E[X_d])$. Furthermore, we assume that data on each stream arrives in blocks of N items. We would like to bound $\text{Pr}_{violation}$, the probability that the arrival of a new block of data

items on a stream will cause a constraint violation at the node monitoring the stream.

We assume that the estimate vector is the expected global statistics vector. Consequently, as long as the local statistics vector held by a node is contained within the distance sphere, the constraint checked by the node is guaranteed not to be violated. That is, if $\|E[\vec{v}(t)] - \vec{v}_i(t)\| < D_{global}$, then $B(E[\vec{v}(t)], \vec{v}_i(t))$ is fully contained in the distance sphere (see Figure 3).

Therefore, the probability that a constraint violation will occur at a node is less than the probability that the local statistics vector held by the node will not be contained in the distance sphere. In other words:

$$Pr_{violation} \leq Pr[\|\vec{v}_i(t) - E[\vec{v}(t)]\| > D_{global}].$$

Since $E[\vec{v}(t)] = E[\vec{v}_i(t)]$:

$$\begin{aligned} Pr_{violation} &\leq Pr[\|\vec{v}_i(t) - E[\vec{v}(t)]\| > D_{global}] \\ &= Pr[\|\vec{v}_i(t) - E[\vec{v}_i(t)]\| > D_{global}] \\ &= Pr[\|\vec{v}_i(t) - E[\vec{v}_i(t)]\|^2 > (D_{global})^2]. \end{aligned}$$

Recall the Markov Inequality, which states that for a random variable Z , that only takes nonnegative values:

$$Pr[Z > \lambda] \leq \frac{E[Z]}{\lambda}.$$

Since $\|\vec{v}_i(t) - E[\vec{v}_i(t)]\|^2$ can be viewed as a random variable taking only nonnegative values, then, according to the Markov Inequality:

$$\begin{aligned} Pr_{violation} &\leq Pr[\|\vec{v}_i(t) - E[\vec{v}_i(t)]\|^2 > (D_{global})^2] \\ &\leq \frac{E[\|\vec{v}_i(t) - E[\vec{v}_i(t)]\|^2]}{(D_{global})^2}. \end{aligned}$$

Let us examine the term $\|\vec{v}_i(t) - E[\vec{v}_i(t)]\|^2$. According to Equation (8):

$$\begin{aligned} \|\vec{v}_i(t) - E[\vec{v}_i(t)]\|^2 &= \left\| \sum_{k=1}^N \frac{\vec{v}_{i,k}}{N} - E[\vec{v}_i(t)] \right\|^2 \\ &= \left\| \sum_{k=1}^N \left(\frac{\vec{v}_{i,k} - E[\vec{v}(t)]}{N} \right) \right\|^2 \\ &= \sum_{j=1}^d \left(\frac{\sum_{k=1}^N (v_{i,k}^{(j)} - E[X_j])}{N} \right)^2. \end{aligned}$$

Since $v_{i,1}^{(j)}, v_{i,2}^{(j)}, \dots, v_{i,N}^{(j)}$ are drawn independently from the random variable X_j ,

$$E[\|\vec{v}_i(t) - E[\vec{v}_i(t)]\|^2] = E \left[\sum_{j=1}^d \left(\frac{\sum_{k=1}^N (v_{i,k}^{(j)} - E[X_j])}{N} \right)^2 \right]$$

$$\begin{aligned}
&= \frac{1}{N^2} \sum_{i=1}^d \sum_{j=1}^N E \left[(v_{i,k}^{(j)} - E[X_j])^2 \right] \\
&= \frac{\sum_{i=1}^d V[X_i]}{N}.
\end{aligned}$$

Therefore:

$$\begin{aligned}
Pr_{violation} &\leq \frac{E[\|\bar{v}_i(t) - E[\bar{v}_i(t)]\|^2]}{(D_{global})^2} \\
&= \frac{\sum_{i=1}^d V[X_i]}{N \cdot (D_{global})^2}.
\end{aligned}$$

If the components of the data vectors are bounded between 0 and 1—as happens in the important case in which they represent probabilities of terms to appear in a document—the Hoeffding bound can be used, to show that.

$$Pr_{violation} \leq \exp \left(-2 \left(D_{global}^2 - \frac{\sum_{i=1}^d V[X_i]}{N} \right)^2 / d \right).$$

Both bounds decrease quickly when D_{global} increases. This suggests that, for data mining applications, features with a small information gain will not cause many constraint violations at any node, since their D_{global} is large. This is practically important, since usually most of the candidate features have a rather small information gain, and thus the proposed algorithm will considerably reduce communication. This is supported by the experimental results presented in the next section.

9. EXPERIMENTAL RESULTS

We performed several experiments with the decentralized algorithm. We tested the algorithm in a distributed feature selection setup. We used the Reuters Corpus (RCV1-v2) [Rose et al. 2002] in order to generate a set of data streams. RCV1-v2 consists of 804414 news stories, produced by Reuters between August 20, 1996, and August 19, 1997. Each news story, which we refer to as a *document*, has been categorized according to its content, and identified by a unique document id.

RCV1-v2 has been processed by Lewis et al. [2004]. Features were extracted from the documents, and indexed. A total of 47,236 features were extracted. Each document is represented as a vector of the features it contains. We refer to these vectors as *feature vectors*. We simulated n streams by arranging the feature vectors in ascending order (according to their document id), and selecting feature vectors for the streams in a round robin fashion.

In the original corpus each document may be labeled as belonging to several categories. The most frequent category documents are labeled with is *CCAT* (the *CORPORATE/INDUSTRIAL* category). In the experiments our goal was to select features that are most relevant to the CCAT category; therefore each vector was labeled as positive if it is categorized as belonging to CCAT, and negative otherwise.

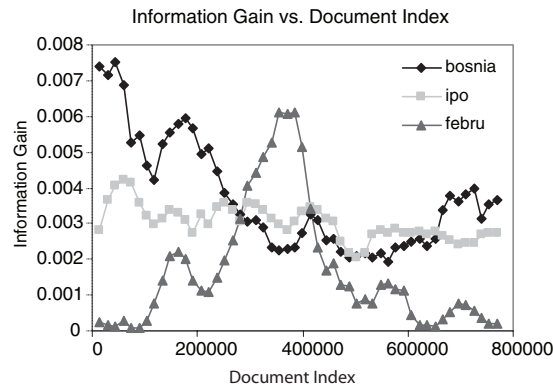


Fig. 4. Information gain for the features *bosnia*, *ipo*, and *febru* as it evolves over the streams. The information gain for the feature *bosnia* displays a declining trend as the stream evolves. The information gain for the feature *ipo* remains relatively steady, while the information gain for the feature *febru* peaks about halfway through the stream.

Unless specified otherwise, each experiment was performed with 10 nodes, where each node held a sliding window containing the last 6700 documents it received. The goal in each experiment was to detect for each feature, at any given time, whether its information gain was above or below a given threshold value. At any given time, the information gain of a feature was based on the documents contained at the time in the sliding windows of all the nodes.

The experiments were designed to explore several properties of the algorithm. We were interested in determining how various parameters of the monitoring task affect the performance of the algorithm. The parameters of the monitoring task can be divided into characteristics of the monitoring task, and tunable parameters. The characteristics of the monitoring task include the number of streams to be monitored, and the desired threshold value. Tunable parameters include the size of the sliding window used by each node, and the permitted error margin. In addition we were interested in examining the behavior of the algorithm when used for simultaneously monitoring several features.

In order to examine the effect of the various parameters on the performance of the algorithm, we chose three features that display different characteristic behavior. The chosen features were *bosnia*, *ipo*, and *febru*. Figure 4 depicts how the information gain for each feature evolved over the streams. The information gain for the feature *bosnia* displays a declining trend as the stream evolve. The information gain for the feature *ipo* remained relatively steady, while the information gain for the feature *febru* peaks about halfway through the stream.

We started by examining the influence of the characteristics of the monitoring task on the performance of the algorithm. Figure 5 shows the number of broadcasts produced when using the decentralized algorithm to monitor each one of the features for threshold values ranging from 0.00025 to 0.006. In addition the cost incurred by the naive algorithm is plotted, that is, the number of messages required for collecting all the data to a central location. One can




Fig. 5. Number of broadcasts produced in order to monitor each feature as a function of the threshold value. In addition, the cost incurred by monitoring a feature by a naive algorithm is plotted. Even for adverse threshold values our algorithm performs significantly better than the naive algorithm.

notice that, even for adverse threshold values, the algorithm incurs a significantly lower communication cost than the cost incurred by the naive algorithm.

In order to check the effect the number of nodes has on the performance of the algorithm, we performed the following experiment: the stream of documents was divided in advance into 100 substreams in a round-robin fashion. Simulations were run with the number of nodes ranging from 10 to 100. In a simulation consisting of n nodes, the first n substreams were used. This methodology ensured that the characteristics of the streams remained similar when simulating different numbers of nodes. Each node held a sliding window of 670 items.

Obviously, increasing the number of nodes will increase the number of broadcasts required in order to perform the monitoring task. Since the nodes in our experiment received streams with similar characteristics, we expected that the number of broadcasts would increase linearly.

Two sets of simulations were run, the first with a threshold value of 0.003, and the second with a threshold value of 0.006. The results are plotted in Figure 6. Both graphs show that the number of broadcasts increases linearly as more nodes are added. Comparing the two graphs reveals that the number of broadcasts increases more moderately when using a threshold value of 0.006. This is due to the fact that as indicated in Figure 4, the average information gain on the monitored features is closer to 0.003.

Next we performed two experiments in order to evaluate the effect of tunable parameters on the performance of the algorithm. We performed the following experiments on the three features: for each feature we chose the threshold value that incurred the highest communication cost (0.0025 for *bosnia*, 0.003 for *ipo*, and 0.00125 for *febru*). We ran a set of simulations on each feature, using error margin values ranging from 0 to 50% of the threshold value. Then we ran an additional set of simulations for each feature, setting the size of

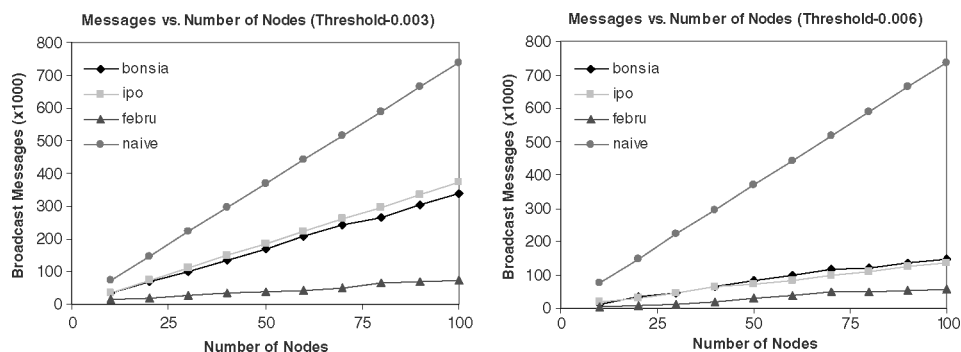


Fig. 6. Number of messages produced in relation to the number of nodes. The number of messages increases linearly as the number of nodes increases, indicating that the algorithm scales well. Since the average information gain of all the features is closer to 0.003 than to 0.006, the number of messages increases more moderately when using a threshold value of 0.006.

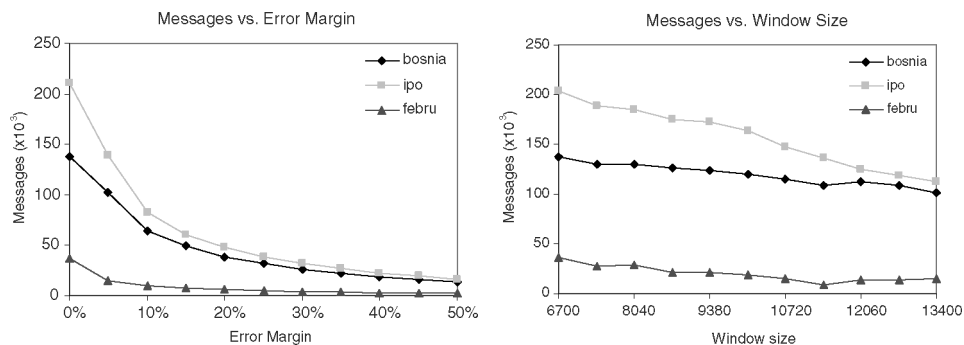


Fig. 7. The influence of tunable parameters on performance. Increasing the error margin is more effective in reducing the communication load than increasing the window size. Using an error margin as small as 5% significantly reduces the communication load.

the sliding window used by each node to values ranging from 6700 items to 13,400 items. The results of these experiments are plotted in Figure 7. The results indicate that increasing the error margin is very effective in reducing the communication load. Using an error margin as small as 5% significantly reduces the communication load. Increasing the window size also reduces the communications load. The effect of increasing the window size is most evident for the feature *ipo*, which incurs the highest communication cost among the three features. In general, increasing the window size has a greater effect the closer the information gain of the feature is to the threshold value.

We performed a set of experiments using the coordinator based algorithm. We ran the coordinator based algorithm on the features using threshold values ranging from 0.00025 to 0.006. We ran three sets of experiments; in the first set no error margin was used. In the second and third sets error margins of 10% and 20% were used. When using no error margin, for certain threshold values the communication cost incurred by the algorithm exceeded the cost incurred by the naive algorithm, while for other threshold values, the algorithm reduced




Fig. 8. Number of messages produced by the coordinator-based algorithm versus threshold value. While the communication cost incurred by the algorithm can exceed the cost incurred by the naive algorithm for certain threshold values, adding a 20% error margin renders the algorithm more efficient for all threshold values.

the communication cost by up to an order of magnitude. Adding a 20% error margin rendered the coordinator-based algorithm more efficient than the naive algorithm for all threshold values. See Figure 8.

Finally, we checked the performance of the algorithm when simultaneously monitoring multiple features using the decentralized algorithm. As the number of features that are monitored simultaneously increases, the probability that a constraint on one of the features will be violated when a new data item is received increases as well. Furthermore, a constraint violation can cause a cascading effect. A constraint violation for a feature at one of the nodes causes all the nodes to calculate a new estimate vector for the feature. Since the value of the estimate vector for the feature has changed, the constraint for the feature may be violated at additional nodes, causing these nodes to broadcast. The purpose of this experiment is to determine the number of simultaneous features the algorithm can monitor while remaining efficient, that is, incurring a cost that is lower than the cost incurred by the naive algorithm. The experiment consisted of a series of simulations, using a threshold value of 0.001. In each simulation a number of features were selected randomly. Simulations were run with the number of features ranging from 1 to 5000. The results of this experiment are plotted in Figure 9. In addition the cost incurred by the naive algorithm is plotted.

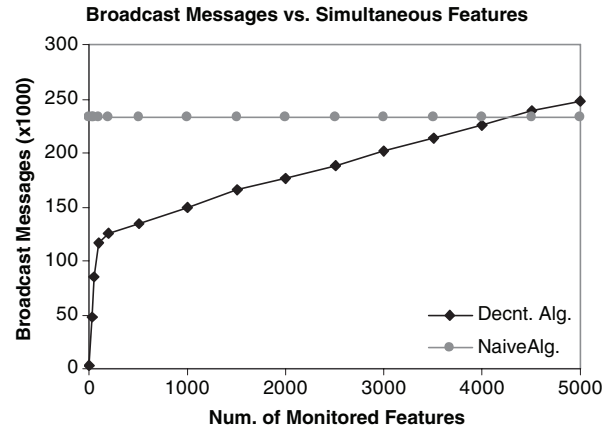


Fig. 9. Number of messages in relation to the number of simultaneously monitored features. Our algorithm remains efficient when simultaneously monitoring up to about 4500 features.




Fig. 10. Number of messages versus threshold value when monitoring the χ^2 statistic.

The results indicate the algorithm remains efficient when simultaneously monitoring several thousands of features, but is inefficient when simultaneously monitoring more than about 4500 features.

An additional function that is commonly used for performing feature selection is the χ^2 statistic. Like information gain, the χ^2 statistic is a function of the global contingency table:

$$Chi(C_{f,X}) = \frac{N(c_{1,1}c_{2,2} - c_{1,2}c_{2,1})^2}{(c_{1,1} + c_{1,2})(c_{1,1} + c_{2,1})(c_{2,2} + c_{1,2})(c_{2,2} + c_{2,1})},$$

where N is the total number of messages the global contingency table represents.

Figure 10 depicts the communication cost incurred by monitoring the χ^2 statistic using the coordinator based algorithm.

10. CONCLUSION

Monitoring streams over distributed systems is an important challenge that has a wide range of applications. Scalability and efficiency of proposed solutions strongly depend on the volume and frequency of communication operations. However, despite the amount of work that has been invested in this direction, most of the efficient solutions found in the literature can only be applied to simple aggregations or to linear functions. Most probably the reason is that when the function is nonlinear, effects seen in one—or only a few—of the streams may often turn out to be misleading with regards to the global picture.

In this work we proposed a solution through a general framework for monitoring arbitrary threshold functions over a system of distributed streams. The evaluation of this approach using real-life data, applied to the information gain function, reveals that it is highly effective in reducing communication frequency.

Immediate future work will concentrate on developing methods to fine-tune various parameters (window size, error margin, threshold) in order to find an optimal tradeoff between communication load and accuracy. We also plan to try and characterize families of functions for which the algorithm is more efficient.

REFERENCES

- ALON, N., MATIAS, Y., AND SZEGEDY, M. 1996. The space complexity of approximating the frequency moments. In *STOC '96: Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. ACM Press, New York, NY, 20–29.
- ARASU, A. AND MANKU, G. S. 2004. Approximate counts and quantiles over sliding windows. In *PODS '04: Proceedings of the Twenty-Third ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. ACM Press, New York, NY, 286–296.
- BABCOCK, B., BABU, S., DATAR, M., MOTWANI, R., AND WIDOM, J. 2002. Models and issues in data stream systems. In *PODS '02: Proceedings of the Twenty-First ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. ACM Press, New York, NY, 1–16.
- BABCOCK, B. AND OLSTON, C. 2003. Distributed top- k monitoring. In *SIGMOD '03: Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 28–39.
- BABU, S. AND WIDOM, J. 2001. Continuous queries over data streams. *SIGMOD Rec.* 30, 3, 109–120.
- BERKOVITZ, L. 2002. *Convexity and Optimization in R^n* . Wiley, New York, NY.
- BULUT, A., SINGH, A. K., AND VITENBERG, R. 2005. Distributed data streams indexing using content-based routing paradigm. In *IPDPS '05: Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium*. IEEE Computer Society Press, Los Alamitos, CA, 94.
- CARNEY, D., ÇETINTEMEL, U., CHERNIACK, M., CONVEY, C., LEE, S., SEIDMAN, G., STONEBRAKER, M., TATBUL, N., AND ZDONIK, S. B. 2002. Monitoring streams—a new class of data management applications. In *VLDB '02: Proceedings of the International Conference on Very Large Data Bases*. 215–226.
- CHARIKAR, M., CHEN, K., AND FARACH-COLTON, M. 2002. Finding frequent items in data streams. In *ICALP '02: Proceedings of the 29th International Colloquium on Automata, Languages and Programming*. Springer-Verlag, London, U.K., 693–703.
- CHERNIACK, M., BALAKRISHNAN, H., BALAZINSKA, M., CARNEY, D., CETINTEMEL, U., XING, Y., AND ZDONIK, S. 2003. Scalable distributed stream processing. In *CIDR 2003: Proceedings of the First Biennial Conference on Innovative Data Systems Research* (Asilomar, CA).
- CORMODE, G., GAROFALAKIS, M., MUTHUKRISHNAN, S., AND RASTOGI, R. 2005. Holistic aggregates in a networked world: Distributed tracking of approximate quantiles. In *SIGMOD '05: Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 25–36.

- DILMAN, M. AND RAZ, D. 2001. Efficient reactive monitoring. In *INFOCOM '01: Proceedings of the Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies*. 1012–1019.
- GIBBONS, P. B. AND MATIAS, Y. 1998. New sampling-based summary statistics for improving approximate query answers. In *SIGMOD '98: Proceedings of the 1998 ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 331–342.
- GIBBONS, P. B. AND TIRTHAPURA, S. 2001. Estimating simple functions on the union of data streams. In *SPAA '01: Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures*. ACM Press, New York, NY, 281–291.
- GIBBONS, P. B. AND TIRTHAPURA, S. 2002. Distributed streams algorithms for sliding windows. In *SPAA '02: Proceedings of the Fourteenth Annual ACM Symposium on Parallel Algorithms and Architectures*. ACM Press, New York, NY, 63–72.
- HUANG, L., GAROFALAKIS, M., HELLERSTEIN, J., JOSEPH, A., AND TAFT, N. 2006. Toward sophisticated detection with distributed triggers. In *MineNet '06: Proceedings of the 2006 SIGCOMM Workshop on Mining Network Data*. ACM Press, New York, NY, 311–316.
- JAIN, A., HELLERSTEIN, J. M., RATNASAMY, S., AND WETHERALL, D. 2004. A wakeup call for Internet monitoring systems: The case for distributed triggers. In *Proceedings of the 3rd ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets, San Diego, CA)*.
- KERALAPURA, R., CORMODE, G., AND RAMAMIRTHAM, J. 2006. Communication-efficient distributed monitoring of thresholded counts. In *SIGMOD '06: Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 289–300.
- LASSERRE, J. B. 2000. Global optimization with polynomials and the problem of moments. *SIAM J. Optimiz.* 11, 3, 796–817.
- LEWIS, D. D., YANG, Y., ROSE, T. G., AND LI, F. 2004. Rcv1: A new benchmark collection for text categorization research. *J. Mach. Learn. Res.* 5, 361–397.
- LIU, L., PU, C., AND TANG, W. 1999. Continual queries for internet scale event-driven information delivery. *IEEE Trans. Knowl. Data Eng.* 11, 4, 610–628.
- MADDEN, S. AND FRANKLIN, M. J. 2002. Fjording the stream: An architecture for queries over streaming sensor data. In *ICDE '02: Proceedings of the 18th International Conference on Data Engineering (ICDE'02)*. IEEE Computer Society Press, Los Alamitos, CA, 555.
- MADDEN, S., SHAH, M., HELLERSTEIN, J. M., AND RAMAN, V. 2002. Continuously adaptive continuous queries over streams. In *SIGMOD '02: Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 49–60.
- MANJHI, A., SHKAPENYUK, V., DHAMDHERE, K., AND OLSTON, C. 2005. Finding (recently) frequent items in distributed data streams. In *ICDE '05: Proceedings of the 21st International Conference on Data Engineering (ICDE'05)*. IEEE Computer Society, Press, Los Alamitos, CA, 767–778.
- MANKU, G. S. AND MOTWANI, R. 2002. Approximate frequency counts over data streams. In *VLDB '02: Proceedings of the International Conference on Very Large Data Bases*. 346–357.
- MOTWANI, R., WIDOM, J., ARASU, A., BABCOCK, B., BABU, S., DATAR, M., MANKU, G., OLSTON, C., ROSENSTEIN, J., AND VARMA, R. 2003. Query processing, resource management, and approximation in a data stream management system. In *Proceedings of the First Biennial Conference on Innovative Data Systems Research (CIDR, Asilomar, CA)*. 245–256.
- OLSTON, C., JIANG, J., AND WIDOM, J. 2003. Adaptive filters for continuous queries over distributed data streams. In *SIGMOD '03: Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 563–574.
- PARRILO, P. 2003. Semidefinite programming relaxations for semialgebraic problems. *Math. Programm.* 96, 2, 293–320.
- ROSE, T., STEVENSON, M., AND WHITEHEAD, M. 2002. The Reuters Corpus Volume 1—from yesterday's news to tomorrow's language resources. In *Proceedings of the Third International Conference on Language Resources and Evaluation (Las Palmas de Gran Canaria)*.
- SHARFMAN, I., SCHUSTER, A., AND KEREN, D. 2006. A geometric approach to monitoring threshold functions over distributed data streams. In *SIGMOD '06: Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 301–312.
- SIDDIQI, K., BOUIX, S., TANNENBAUM, A., AND ZUCKER, S. W. 2002. Hamilton-Jacobi skeletons. *Int. J. Comput. Vis.* 48, 3 (July), 215–231.

- SULLIVAN, S., SANDFORD, L., AND PONCE, J. 1994. Using geometric distance fits for 3-d object modeling and recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 16, 12, 1183–1196.
- TERRY, D., GOLDBERG, D., NICHOLS, D., AND OKI, B. 1992. Continuous queries over append-only databases. In *SIGMOD '92: Proceedings of the 1992 ACM SIGMOD International Conference on Management of Data*. ACM Press, New York, NY, 321–330.
- VARADHAN, G., KRISHNAN, S., SRIRAM, T., AND MANOCHA, D. 2004. Topology preserving surface extraction using adaptive subdivision. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*. ACM Press, New York, NY, 235–244.
- YI, B.-K., SIDIROPOULOS, N., JOHNSON, T., JAGADISH, H. V., FALOUTSOS, C., AND BILIRIS, A. 2000. Online data mining for co-evolving time sequences. In *ICDE '00: Proceedings of the 16th International Conference on Data Engineering*. IEEE Computer Society, Press, Los Alamitos, CA, 13.
- ZHU, Y. AND SHASHA, D. 2002. Statstream: Statistical monitoring of thousands of data streams in real time. In *VLDB '02: Proceedings of the International Conference on Very Large Data Bases*. 358–369.

Received September 2006; revised May 2007; accepted May 2007