Approximate Decision Making in Large-Scale Distributed Systems

Ling Huang[†] Minos Garofalakis[◇] Anthony D. Joseph[*] Nina Taft[†]
[†]Intel Research [◇] Yahoo! Research [*] UC Berkeley

1 Introduction

As the Internet has evolved into a valuable and critical service platform for business and daily life, the research community has enthusiastically applied data mining methods to improve application performance by analyzing and optimizing the behaviors of the underlying systems (*e.g.,* datacenter design, network resource provisioning, network security, etc.) These data mining procedures often use large-scale widely-distributed monitoring systems, which continuously generate numerous distributed data streams, and backhaul *all of the data* to a central location (*e.g.,* a Network Operation Center or NOC) for data analysis and decision making. This application scenario presents both new opportunities and challenges in efficient data analysis and online decision making, where a decision function depends on aggregating and analyzing *continuous* data streams from *distributed* monitors.

The statistics and machine learning communities have performed extensive research into decision making methods [1], including outlier detection, clustering, classification, etc., with the results being algorithms that mainly assume all data have been collected at a central point, and focus on post-collection data analysis and problem diagnosis, with little consideration of the more general distributed, continuous data collection and analysis problem. We believe that the machine learning community should now focus on the design of algorithms that function well with limited data.

We envision two open problems: efficiently performing online decision making *with low communication overhead*, and providing fine-grain control over the tradeoff between decision accuracy and communication overhead. Most existing research has focused on sampling techniques, however, the randomness in this type of sampling could discard key information needed by decision making algorithms. Instead, we advocate using *smart filtering* for data reduction, where the filtering is designed to carefully select which data to *not* ship. Specifically, the filtered data should be that which has minimal impact on decision making performance or its accuracy.

2 Challenges

Using a centralized model, where all monitored data is periodically pushed to the NOC, simplifies the application of detection and correlation functions to global data. However, a centralized model introduces significant efficiency, timescale, and size scalability limitations, especially when attempting to collect monitoring data and perform detection on sub-second or smaller time scales; and increase the number of monitors by an order of magnitude or more. Changes in timescale and size can massively increase the volume of data sent to the NOC, potentially overloading its network capacity.

The distributed nature of monitor sites implies important *communication network constraints* due to either network bandwidth restrictions or power limitations (*e.g.,* sensor battery life). Such limitations are obvious in sensor networks, but also in large enterprise networks which typically do not over provision links to remote office sites. Clearly, we need *communication-efficient distributed monitoring* because naïve solutions that simply continuously "push" complete data streams to a central collection site will not scale.

The database community has developed approximate data replication protocols for managing distributed and continuous data streams [5], that efficiently and effectively enable centralized access to distributed streams (data objects whose values continuously change over time). Because their goal is to aggregate data at the NOC within an ϵ-error bound on accuracy *regardless of actual system conditions*, stream processing approaches suffer from excessive query overhead in the presence of bursty data and are ill suited for efficient decision making.

Our key insight in the efficient decision making problem is that *exact data is often not a requirement* — the important metric is not ϵ-error approximation of system state, rather it is ϵ-error decision making. In other words, we care about how accurately monitoring detects a violation, not how accurately it determines overall system state.

With this insight we can use approximate replication techniques to reduce communication costs and study fundamental tradeoffs between central site synchronization communication costs and decision making accuracy. Using machine learning and statistics, this *cost-accuracy tradeoff* can be codified in a system that lets users specify a minimum allowable accuracy level, and then it minimizes communication costs while meeting the specified accuracy requirement.

3 The Problem Space

In Figure 1, we show the three axes of the design space for the online decision making and detection problem:

Time Scales of detection represent at least three detection condition types. *Instantaneous* triggers fire when an aggregate threshold value is violated at any single instant [4, 6], while *fixed-window* and *cumulative* triggers detect persistent threshold violations over fixed and *any size* windows of time, respectively [2].

Detection Functions include simple linear (*e.g.,* `SUM`) and more sophisticated ones (*e.g.,* Top-k, PCA, and SVM) enabling a wide range of detection applications (*e.g.,* botnet attacks, volume anomalies in ISP network, and electric power grid anomalies).

Communication Architectures include a one-level tree where every monitor directly communicates with the NOC, multi-level tree structures where monitors have a parent-child relationship, and pure distributed topologies.

Our research has explored a slice of the problem space: various detection functions defined on a one-level tree topology (the shaded area in Figure 1). The rest of the space represents an interesting opportunity for the research community.

4 Towards Efficient Decision Making

As an exploratory step, we propose, D-Trigger, a general framework for efficient online detection that gracefully integrates a variety of decision making, online machine learning, approximation, and optimization algorithms. Our key goals and accomplishments are to: enable real-time detection where a system's state is tracked continuously, so even small abnormal events are detected; significantly reduce the data collected for detection, thus reducing communication overhead; guarantee desired detection accuracy even with a reduced amount of data. D-Trigger combines very high detection accuracy and low communication overhead for the detection of various unusual events (*e.g.,* detecting botnet attacks, network traffic volume anomalies, and electric power grid anomalies).

We have developed two specific approaches for security applications: a queueing-based approach for botnet detection [2] and a Principal Components Analysis-based approach for network-wide anomaly detection [3]. A common theme in both approaches is collaborative anomaly detection across many widely distributed monitors, and a key lesson we have learned is that data can be intelligently filtered by controlling or bounding detection errors. In our approaches, distributed monitors perform local information processing and only send approximate (filtered) data to the NOC. Because the NOC has imperfect knowledge of the monitored data, it may make mistakes in the detection, but we leverage machine learning and statistics to design monitor and NOC protocols that bound detection error (providing high accuracy), while simultaneously minimizing communication — yielding the desired fine-grained tradeoff between detection accuracy and communication cost.

Our research shows that often 80 to 90% of monitor data can be filtered – a result of our focus on accurate anomaly detection, instead of accurate state approximation. In a well-performing system (the common case), anomalies are rare events and the NOC does not need an accurate view of system state. It is only when system state approaches the decision making threshold that the NOC needs a more accurate view (*i.e.,* more detailed data from monitors). This is where we apply machine learning and statistics techniques to provide monitors with dynamic feedback for adaptive data filtering.

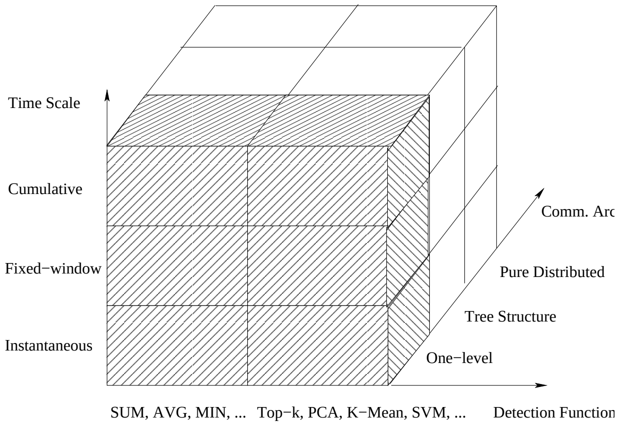

Figure 1: The whole problem space.

Providing efficient decision making in large-scale distributed systems remains an open problem, however our initial results from leveraging machine learning demonstrate the feasibility of balancing accurate decision making with minimized communication needs. Based on the applications we have examined so far, we believe that our framework and approach are broadly applicable and a basis for exploring a wide spectrum of algorithms that deal with anomaly detection. There are several research directions for further exploration, including using multi-level tree or pure distributed communication architectures to further reduce the processing and communication burden at the NOC; supporting more sophisticated types of detection algorithms (*e.g.,* wavelet decomposition, entropy analysis, clustering, classification, and sequential hypothesis methods); and developing resilient monitoring infrastructures that can tolerate data losses.

References

[1] HASTIE, T., TIBSHIRANI, R., AND FRIEDMAN, J. *The Elements of Statistical Learning.* Springer, 2001.

[2] HUANG, L., GAROFALAKIS, M., JOSEPH, A. D., AND TAFT, N. Communication-efficient tracking of distributed cumulative triggers. In *Proceedings of International Conference on Distributed Computing Systems (ICDCS)* (2007).

[3] HUANG, L., NGUYEN, X., GAROFALAKIS, M., HELLERSTEIN, J. M., JOSEPH, A. D., JORDAN, M. I., AND TAFT, N. Communication-efficient online detection of network-wide anomalies. In *Proceedings of 26th Annual IEEE Conference on Computer Communications (INFOCOM)* (2007).

[4] KERALAPURA, R., CORMODE, G., AND RAMAMIRTHAM, J. Communication-efficient distributed monitoring of thresholded counts. In *Proceedings of ACM SIGMOD* (2006).

[5] OLSTON, C. *Approximate Replication.* PhD thesis, Stanford University, June 2003.

[6] SHARFMAN, I., SCHUSTER, A., AND KEREN, D. A geometric approach to monitoring threshold functions over distributed data streams. In *Proceedings of ACM SIGMOD* (2006).