

Network-wide Anomaly Event Detection and Diagnosis with perfSONAR

Yuanxun Zhang, Saptarshi Debroy, Prasad Calyam
University of Missouri-Columbia

Email: yzd3b@mail.missouri.edu, {debroy, calyamp}@missouri.edu

Abstract—High Performance Computing (HPC) environments supporting data-intensive applications need multi-domain network performance measurements from open frameworks such as perfSONAR. Detected network-wide correlated anomaly events that impact data throughput performance need to be quickly and accurately notified along with a root-cause analysis for remediation. In this paper, we present a novel network anomaly events detection and diagnosis scheme for network-wide visibility that improves accuracy of root-cause analysis. We address analysis limitations in cases where there is absence of complete network topology information, and when measurement probes are miscalibrated leading to erroneous diagnosis. Our proposed scheme fuses perfSONAR time-series path measurements data from multiple domains using Principal Component Analysis (PCA) to transform data for accurate correlated and uncorrelated anomaly events detection. We quantify the certainty of such detection using a measurement data sanity checking that involves: (a) measurement data reputation analysis to qualify the measurement samples, and (b) filter framework to prune potentially misleading samples. Lastly, using actual perfSONAR one-way delay measurement traces, we show our proposed scheme’s effectiveness in diagnosing the root-cause of critical network performance anomaly events.

Keywords-Multi-domain Network Performance Monitoring, Anomaly Event Detection, Root-cause Diagnosis Certainty

I. INTRODUCTION

Distributed computing applications are increasingly being developed in scientific communities in areas such as biology, geography and high-energy physics. These communities transfer data on a regular basis between computing and collaborator sites at high-speeds on multi-domain networks that span across continents. To ensure high data throughputs through effective network monitoring, there is a rapidly increasing trend to deploy multi-domain, open measurement frameworks such as perfSONAR [1]. The perfSONAR framework has been developed over the span of several years by worldwide-teams and has over 1400 measurement points all over the world.

However, providing scientists and network operators with a network-wide performance visibility based on the perfSONAR measurement archives within data-intensive science collaborations such as [2] poses several challenges [3–5]. It requires automated techniques to query, analyze, detect and diagnose prominent network performance anomaly events that hinder data transfer performance. The general lack of network topology information accompanying the multi-domain measurements data compounds the challenges in root-cause diagnosis of performance bottlenecks. More specifically, it is non-trivial to identify and locate network-wide anomaly events

This material is based upon work supported by the US Department of Energy under Award Numbers: DE-SC0001331 and DE-SC0007531. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof.




Fig. 1: perfSONAR dashboard with throughput measurement notifications for ESnet to European sites

that impact data throughput performance without publicly accessible topology services for measurement points [6, 7]. The identification and location diagnosis of anomaly events is particularly challenging in cases with measurement data spanning multiple network paths.

Fig. 1 for example shows a typical perfSONAR dashboard with color-coded periodic throughput (‘Reds’ are ≤ 100 Mbps, ‘Yellows’ are < 500 Mbps, and ‘Greens’ are ≥ 500 Mbps) measurement status event notifications for different paths, specifically between ESnet and several European sites. Although the dashboard serves the purpose of interesting events notification, pertinent issues essential to ascertain the significance of such events remain unanswered, such as: Do the events in Sets II and III correspond to a common network anomaly event? If yes, then what are the root-causes of such anomaly events? Do events in Set III belonging to the same destination signify anomaly correlation? Do events in Set II belonging to the same source signify anomaly correlation?

Answering such critical questions for effective troubleshooting becomes even more challenging as the publicly accessible measurement samples collected from perfSONAR deployments often have measurement mis-calibration or issues such as invalid measurement data. Examples of issues include negative one-way delay values due to faulty clock synchronization between measurement servers. Such issues result in erroneous features [8], or too dense/sparse or irregular (i.e., long data collection gaps) measurement sampling frequency leading to missed anomaly events and exponential anomaly detection time [9]. Such measurement mis-calibration eventually manifests in triggering of erroneous detections and useless diagnosis/notifications.

In this paper, we present a *novel scheme that can fuse time-series of perfSONAR path measurements from multiple domains with common intermediate hops for: (a) correlated anomaly event detection, and (b) a simultaneous sampling*

trend analysis for accurate and timely notifications. The anomaly event detection involves fusion of multiple time-series to transform perfSONAR measurements onto new axes through PCA [10] (i.e., principal component analysis), which obviates the need of complete network topology information. This transformation extracts common features upon which our earlier adaptive plateau event detection (APD) scheme [9] is applied to detect uncorrelated anomaly events (change-points from statistical norm) at a network-wide level. These detected network anomalies are then compared against the Q-statistic threshold to isolate the correlated anomaly events. This approach leverages the fact that the PCA technique is best suited to be configured by network operators as a “black-box” [11, 12] for correlation analysis.

To address the problem of misleading data use within the anomaly detection analysis, we perform measurement data sanity checking using an adaptive measurement data *Reputation Analysis* coupled with a novel *Filter Framework*. The *Reputation Analysis* scheme assigns reputation scores to measurement paths based on domain-specific historical sampling trends involving factors such as: validity of the measured data [8], and sampling periodicity [9], that may potentially cause measurement mis-calibration. The multi-path reputation scores are then translated into “certainty” of detection quantification which provides a network-wide meta-perspective for the network operators in a multi-domain environment. The *Filter Framework* is used subsequently to apply *temporal and spatial filters* to the multi-path measurements for root-cause identification of uncorrelated anomalies, as well as for pruning the misleading measurement features in case of “low certainty of detection” of anomaly events.

Using synthetic data mimicking actual perfSONAR traces, we compare our scheme with similar anomaly detection schemes demonstrating the scheme’s effectiveness with high detection accuracy and low false positive rate. We also implement the proposed data sanity checking scheme in our Narada Metrics framework [13] that features several perfSONAR extensions, and is being used in actual multi-domain enterprises. We use the Narada Metrics framework to collect both short-term (one day) and long-term (one month) perfSONAR one-way delay measurement datasets from United States Department of Energy (DOE) lab sites (e.g., FNAL, ORNL) and perform multiple case studies for performance evaluation of the proposed scheme. Using these case studies, we demonstrate that our scheme can fuse multi-domain measurement data in order to: (a) effectively ascertain correlation among anomaly events, (b) leverage a source-side vantage point to diagnose whether an anomaly event location is local or in an external domain, (c) pin-point potential root-cause locations for both correlated and uncorrelated anomalies without complete network topology information, and (d) intelligently prune potentially misleading features in the measurement data to increase the certainty of detection.

The remainder paper organization is as follows: Section II describes the related work. Section III presents background on plateau detection and the PCA technique. Section IV presents our PCA-APD-Q-statistic scheme. Section V discusses the certainty quantification of anomaly events. In Section VI, we evaluate the accuracy of our proposed scheme and perform case studies to isolate bottleneck anomaly event locations with actual measurement traces. Section VII concludes the paper.

II. RELATED WORK

A. Network Anomaly Event Detection Techniques/Tools

To assist network operators in troubleshooting bottlenecks (e.g., prolonged congestion events or device mis-configurations) in multi-domain high-speed networks, a number of smart and effective network monitoring tools based on statistical measurement data analysis techniques, such as [14–21] have been developed. Particularly, in [14], the authors provide a user-level Internet diagnosis tool which is used for diagnosing network performance problems. A passive network monitoring system is described in [15] that monitors traffic between PlanetLab sites to detect anomalous behavior. Further, in [16], the authors propose Information Plane (iPlane), designed as a service to obtain information about Internet conditions. Authors in [17] present Crowdsourcing Event Monitoring (CEM) approach to detect, isolate and report service-level network events. *Many of these analysis techniques/tools however lack automation provided by our work, and are not useful in perfSONAR measurement data context to ascertain anomaly event correlation for network-wide performance visibility and effective troubleshooting.*

B. Topology-dependent Correlated Anomaly Detection

Alternately, there have been works such as [3, 6, 7, 12, 22–25] that use network topology information for correlated anomaly event detection to localize bottlenecks. Authors in [12] use Kalman-filter for anomaly detection and build a traffic matrix of an enterprise network to overcome link basis limitations. A root-cause analysis and anomaly localization tool called Pythia is proposed in [3] that uses perfSONAR measurements. In [25], a Service-quality Characterization of Internet-path (SCI) scheme is proposed that relies on delay and loss measurements collected from vantage points at two ends of a path. Similarly in [26], use of QoS parameters collected from vantage points at two ends of network paths for detecting network anomaly events is proposed. Our work closely relates to NICE (Network-wide Information Correlation and Exploration) framework proposed in [6] for analyzing anomaly events through data correlations. In our recent work [7], we used topology-aware anomaly detection for location diagnosis of correlated anomaly events. Most of these prior works have a strict requirement for complete topology information, which is a well known open research problem as discussed in [27]. Whereas, in this work, we propose a partially topology-agnostic network-wide anomaly event detection and diagnosis scheme for perfSONAR deployments. *Our work’s novelty is that we address anomaly event cases lacking publicly available topology information accompanying measurement data sets to isolate bottleneck root-cause location.*

C. PCA-based Correlated Anomaly Detection

PCA based measurement data projection schemes, such as [11, 28–30] have recently been proposed by researchers to detect and diagnose anomalies in the absence of network topology information. Authors in [11] use PCA technique on passive measurements for network anomaly detection on a network link basis. A PCA subspace projection methodology is proposed in [29] where the authors apply PCA on data that have already undergone random projection to detect anomaly events. In our earlier work [30], we used PCA to isolate and diagnose the locations of the correlated anomalies in the network in the absence of complete network topology information. *Our*

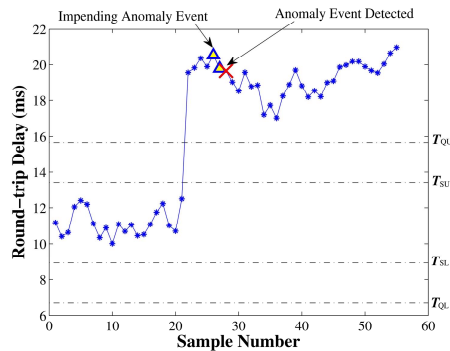


Fig. 2: Plateau-detector thresholds illustration

work builds upon these earlier works, and extends them in context of measurement data reputation analysis and filtering to address cases where misleading data in the measurement samples collected from perfSONAR archives impact anomaly detection accuracy.

D. Measurement Data Sanity Checking

Guidelines for measurement best practices and the perils of using potentially misleading data were first outlined in [8]. Our work on using sanitized measurement data for anomaly detection is closest to the work by authors in [31], where an anomaly detection system is developed based on prediction of upper and lower dynamic thresholds of various time-varying data trends. Reputation-based trust schemes have long been used by the scientific community for decision making in shared environments. Feedback-based reputation management schemes have been proposed for large open environments in e-commerce [33], peer to peer (P2P) computing [34], and wireless systems [35]. *Our work is the first in effectively using reputation-based sanitized measurement data gathering, and lays the foundation for addressing the reputation management of measurement points or domains.*

III. BACKGROUND

In this section, we first define anomaly events that are of interest to network operators, and give an overview of adaptive plateau event detection (APD) that we rely in this paper for automatic anomaly event notifications. Following this, we formally introduce the PCA technique which we will leverage along with APD and Q-statistic in order to establish correlation between such network-wide anomaly events.

A. Anomaly Events

One of the significant challenges in dealing with perfSONAR measurement datasets is to decide which kind of network events (i.e., ‘Reds’ in Fig. 1) need to be labeled and notified as anomaly events that may affect data-intensive application performance bottlenecks. Various traffic related anomaly events are caused due to IP route/AS path change events that involve traffic re-routing on backup paths due to ISP traffic migration for maintenance reasons. These events manifest in the form of spikes, dips, bursts, persistent variations and plateau trends in network performance metrics such as round-trip delay, available bandwidth and packet loss obtained through end-to-end active measurements. Based on documented experiences from network operators and data intensive science application users [4] and based on our own

discussions with other HPC network operators (e.g., ESnet, Internet2, GEANT), the notification of ‘plateau anomalies’ shown in Fig. 2 are the most worthy to be notified. These anomaly events are commonly known to impact data transfer speeds at the application-level on high-speed network paths.

B. Adaptive Plateau Detector

Network operators, when analyzing a measurement time-series of network performance metrics, typically look for plateau event trends through visual inspections and seek for automated notification of such network-wide detected anomaly events. Variants of plateau anomaly event detectors have been developed and adopted in large scale monitoring infrastructures such as NLANR AMP [36] and SLAC IEPM-BW [4], which are predecessors to the perfSONAR deployments. These detectors use static configurations of ‘sensitivity’ and ‘trigger elevation threshold’ parameters to detect that a plateau event or a ‘change event’ has occurred.

A plateau event is detected if the most recent measurement sample value crosses the upper or lower thresholds of the summary (i.e., T_{SU} , T_{SL}) and quarantine (i.e., T_{QU} , T_{QL}) buffers as determined by the settings of sensitivity and trigger elevation parameters. The summary buffer is used to maintain sample history that indicates the normal state (before anomaly event occurs), and a quarantine buffer is used to store outlier data samples that are twice the normal state sample values. The sample counts in above buffers are used to maintain trigger count values over a pre-configured trigger duration before an alarm of anomaly event occurrence (indicated by the cross mark in Fig. 2) is notified. The trigger duration before samples are marked for impending anomaly states (triangle symbols shown in Fig. 2) should be chosen long enough to avoid false alarms due to noise events corresponding to intermittent spikes, dips, or bursts.

Our earlier adaptive plateau-detector (APD) algorithm [9] scheme avoids manual calibration of ‘sensitivity’ and ‘trigger elevation threshold’ parameters and has been shown to be more accurate than earlier static plateau detection schemes [36] [4] over diverse profiles of measurement samples on network paths. Given that we rely on APD in this paper for detecting anomaly events in the measurements, we will illustrate the advantages of using APD later in Section VI-A. Specifically, we show how APD outperforms plateau detectors using static thresholds (SPD) in detecting uncorrelated anomalies of smaller magnitudes with much fewer false alarms.

C. Principal Component Analysis and Q-statistic

Root-cause analysis of detected anomaly events at a network-wide level in the absence of complete network topology information is non-trivial as explained in Section I. Establishing correlation between anomaly events is important to not only diagnose the detected anomaly event location, but also to determine whether resolving cause of one event can auto-resolve multiple other related events in the troubleshooting process.

We use PCA technique along with a Q-statistic [10, 37] test on perfSONAR multi-path time series data in order to isolate correlated anomaly events. The reason to use PCA is because it is a dimensionality-reduction approach that involves mapping a set of data points within time-series onto new coordinates. The new coordinates are called the principal axes or principal components that help to extract common features in the data

points of multiple time-series, and thus visually separate the normal behavior from anomalous behavior.

Let \mathbf{Y} be the $n \times m$ time-series measurement matrix, which denotes the time-series of all links and centered to have zero mean, with n being the number of rows and m being the number of columns. Thus, each column denotes the time-series of the i -th link and each row j represents an instance of all the links. Applying PCA to \mathbf{Y} yields a set of m principal components, $\{\mathbf{v}_i\}_{i=1}^m$, where the first principal vector \mathbf{v}_1 is given as:

$$\mathbf{v}_1 = \arg \max_{\|\mathbf{v}\|=1} \|\mathbf{Y}\mathbf{v}\| \quad (1)$$

Where $\|\mathbf{Y}\mathbf{v}\|$ is proportional to the variance of the data measured along \mathbf{v} . Proceeding iteratively, the k -th principal component \mathbf{v}_k is given as:

$$\mathbf{v}_k = \arg \max_{\|\mathbf{v}\|=1} \left\| \left(\mathbf{Y} - \sum_{i=1}^{k-1} \mathbf{Y}\mathbf{v}_i\mathbf{v}_i^T \right) \mathbf{v} \right\| \quad (2)$$

The first principal component \mathbf{v}_1 captures the maximum variance. The next principal component captures the maximum variance among the remaining orthogonal directions. After choosing the principal components or axes, the dataset can be projected onto the new axes. The subspace method that we use separates principal components into normal and abnormal principal components. The normal principal components reside in the normal subspace S_{no} whereas the abnormal principal components reside in the abnormal subspace S_{ab} . In the pioneering PCA work [11], the authors observed that the normal measurements, i.e., lower k components reside in S_{no} , and the abnormal measurements i.e., $(n-k)$ components reside in S_{ab} . From our analysis of large number of perfSONAR measurement traces, we found that the correlated anomaly events always reside in the lower k components or S_{no} , subspace and uncorrelated anomaly events always reside in the $(n-k)$ components or S_{ab} subspace. This finding of ours is consistent with the findings in [11] in terms of establishing correlation among measurement traces. Hence, the observation from our experiments is guiding our decision to use the subspace method to effectively separate correlated and uncorrelated anomaly events by selecting the lower k components.

Now let $\mathbf{y} = \mathbf{y}(t)$ denote a n -dimensional vector of measurements (for all links) from a single time step t . Detection of anomalies relies on the decomposition of link measurements $\mathbf{y} = \mathbf{y}(t)$ at any step into normal and abnormal components, $\mathbf{y} = \mathbf{y}_{no} + \mathbf{y}_{ab}$, the \mathbf{y}_{no} corresponds to modeled normal measurements (the projections of \mathbf{y} onto S_{no}), and the \mathbf{y}_{ab} corresponds to residual measurements (the projections of \mathbf{y} onto S_{ab}), and can be computed as:

$$\begin{aligned} \mathbf{y}_{no} &= \mathbf{P}\mathbf{P}^T\mathbf{y} = \mathbf{C}_{no}\mathbf{y} \\ \mathbf{y}_{ab} &= (\mathbf{I} - \mathbf{P}\mathbf{P}^T)\mathbf{y} = \mathbf{C}_{ab}\mathbf{y} \end{aligned} \quad (3)$$

where $\mathbf{P} = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_k]$ is formed by the first k principal components which capture the dominant variance in data. The matrix $\mathbf{C}_{no} = \mathbf{P}\mathbf{P}^T$ represents the linear operator that performs projection onto normal subspace S_{no} , and the \mathbf{C}_{ab} represents the projection onto the abnormal subspace S_{ab} .

As described in [11], a volume anomaly event typically results in a large change to \mathbf{y}_{ab} ; thus, a useful metric for detecting abnormal measurements pattern is squared prediction error (SPE):

$$\text{SPE} \equiv \|\mathbf{y}_{ab}\|^2 = \|\mathbf{C}_{ab}\mathbf{y}\|^2 \quad (4)$$

We consider network measurements to be normal if $\text{SPE} \leq \delta^2$, where δ^2 denotes the threshold for the SPE at the $1 - \alpha$ confidence level. Such a statistic test for the SPE residual function is known as Q-statistic, which was developed in [10] to deal with residuals related to principal component analysis. We use the Q-statistic to analyze the significance of the differences among the data sets by capturing the correlated anomalies in the first k principal components that reside in the *normal* measurements subspace S_{no} .

IV. ANOMALY EVENT DETECTION

In this section, we present our proposed PCA-APD-Q-statistic based network anomaly event detection and diagnosis.

A. Scheme Overview

In Fig. 3, we show the components and steps involved in our proposed network-wide anomaly event detection and certainty diagnosis scheme. The steps begin with data collection through querying of distributed measurement archives (accessible at an address e.g., [curl http://fnal-owamp.es.net:8085/esmond/perfsonar/archive](http://fnal-owamp.es.net:8085/esmond/perfsonar/archive)) by using perfSONAR-compliant web service clients. The site list of measurement archives (MAs) that are available for query can be selected using a global lookup service hosted by the perfSONAR community. This service registers the addresses of all openly-accessible measurement archives within individual domains. Upon data collection, the multi-path time-series data is fed simultaneously to Anomaly Event Detection and Data Sanity Checking components.

The Data Sanity Checking component performs reputation analysis over the collected samples. In case of a correlated anomaly event detection, the root-cause location isolation is easier, and the reputation of the entire measurement data set influences the overall certainty of the detection. However, uncorrelated anomaly events are rather difficult to isolate in the absence of network topology information as explained in Section I, and thus requires further processing that involves passing the multi-path time-series data through the filter framework. Output of the filter framework dictates whether to prune the potentially misleading data or perform a destination-to-source conversion to look for correlated anomaly events. In either case, recursive anomaly detection is performed until correlated anomaly events are detected and root-cause locations are isolated. Finally, Anomaly Event Detection and Data Sanity Checking schemes together enable the ranking of detected events by certainty of detection in order to provide a network-wide meta-perspective for effective troubleshooting.

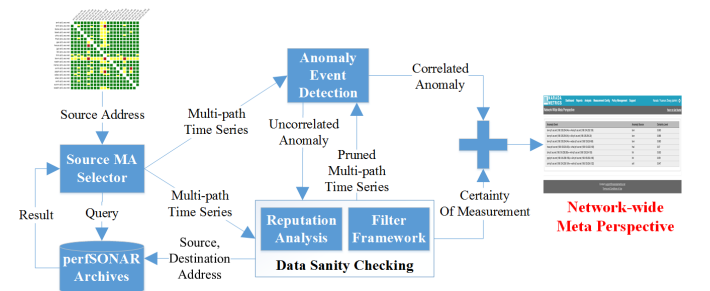


Fig. 3: Schematic diagram of our proposed network-wide anomaly event detection and certainty diagnosis

B. PCA-APD-Q-statistic Analysis

Fig. 4 shows the sequence of steps involved in our PCA-APD-Q-statistic based anomaly detection and diagnosis scheme. Through standardized request/response messages, active measurement time series data relating to end-to-end performance measurement tools such as OWAMP (one-way delay as specified in IETF RFC 4656) are downloaded for any given site (i.e., Source Site A). The downloaded multi-path time series datasets are in the form of JSON object files, which are then processed using parsing for applying PCA technique in the subsequent step.

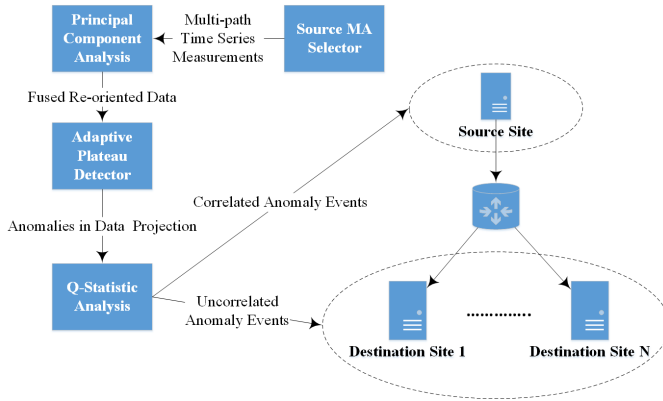


Fig. 4: PCA-APD scheme components diagram

The output of PCA is the fused reoriented data comprising of eigen vectors, where the first eigen vector captures maximum variability and the last is left with minimum variability. What this translates into in-reality is that - the data projection using the first eigen vector has variability that is common to most of datasets and the last eigen vectors have the variability that is least common in the dataset (e.g., variability present in only one dataset amongst all). Next, data dimension selection is performed on the fused reoriented data. For example, if we are interested only in the common anomalies, we will select only the first principal component as described in the previous section. After the data dimension (number of eigen vectors) is selected, the data is projected using the principal components, and is passed as input for APD algorithm to detect anomalies.

Although most of the correlated anomalies subspace are captured in the first, or first and second principal components, it is likely that the normal subspace is also located in the lower k -components. In order to accurately capture all of the anomaly events within measurement time-series, we leverage our APD scheme on the PCA transformed (or fused reoriented) measurement data. To further classify the correlated and uncorrelated anomaly events, we employ the Q-statistic test described earlier in Section III-C. Moreover, if we find that the site-of-interest (i.e., Source Site A) is featured in many or all of the correlated anomaly event paths, we can conclude that the anomaly event root-cause is local. If otherwise, we can conclude that the anomaly event root-cause is in an external domain, and above sequence of diagnosis steps can be applied to other domains whose measurement data is accessible with the hope of localizing the root-cause in one of the external domains.

To substantiate the above rationale for correlated and uncorrelated anomalies, we use synthetic time-series measurements for study purposes that comprise of 16 traces of one-way

delay measurements collected from perfSONAR archives that do not have any anomaly events. Into these traces, we inject 5 anomaly events within a common time period window to create a correlated anomaly event, and also inject 16 uncorrelated anomaly events in other time period windows.

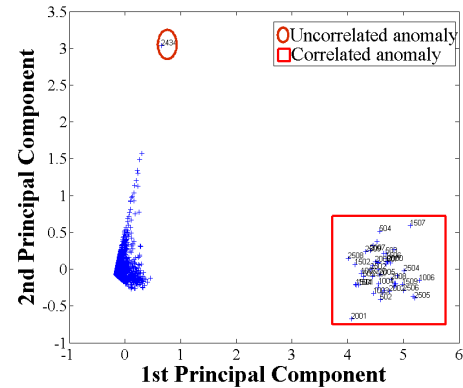


Fig. 5: Correlated and uncorrelated anomaly subspace separation with PCA application

As shown in Fig. 5, all the correlated anomaly events are captured in the first principal component, and an uncorrelated anomaly event is captured in the second principal component. In repeated studies with different synthetic measurement time-series, we found that all the correlated anomaly events are captured mostly in the first principal component, and at worst in the second principal component in a very few number of instances.

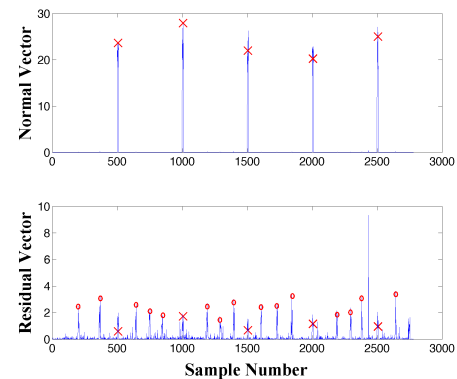


Fig. 6: Measurements of normal space vector squared magnitude ($\|y_{no}\|^2$, upper), and residual space vector squared magnitude ($\|y_{ab}\|^2$, lower) for the synthetic data

As shown in Fig. 6, we separate the link measurements y into normal subspace and residual subspace. The lower part of the figure shows the SPE of y 's projection in the residual subspace y_{ab} , and the upper part shows y 's projection in the normal subspace y_{no} . On these plots, we have marked the correlated anomalies with crosses (x) and uncorrelated anomalies with circles (o). In the lower part of the figure, it is clear that the magnitude of the residual vector y_{ab} is dominated by uncorrelated anomalies rather than correlated anomalies. As a result, it is difficult to discern the correlated and uncorrelated anomalies in the residual vector y_{ab} . However, in the upper part of the figure, only correlated anomalies along with normal

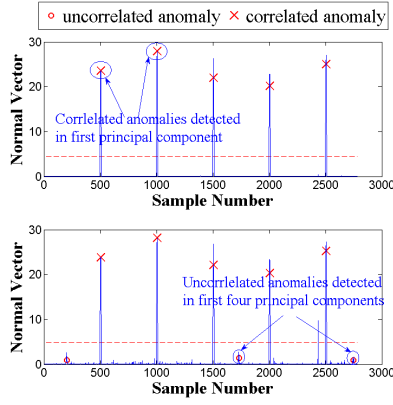


Fig. 7: Measurements of normal space vector squared magnitude $\|\mathbf{y}_{no}\|^2$ ($\mathbf{P} = [\mathbf{v}_1]$, upper) and ($\mathbf{P} = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4]$, lower) for the synthetic data

measurement data are captured in the projection. Thus, the magnitude of normal measurement data is obviously different from the correlated anomaly measurement data, which makes the detection of anomalies much easier to distinguish.

Above observation shows that the normal vector \mathbf{y}_{no} is suitable to detect correlated anomalies at a network-wide level. However, we still want to find uncorrelated anomalies. In Fig. 6, only correlated anomalies are captured in the normal vector. Although residual vector can capture all the correlated and uncorrelated anomalies, it is difficult to discern them because only the first principal axis is selected in the Fig. 6 to capture normal traffic and correlated anomalies. Hence, we need to increase principal axes to capture uncorrelated anomalies.

In the top-portion of Fig. 7, only correlated anomalies are captured in the first principal component projection. However, in the lower plot of Fig. 7, correlated anomalies and some of the uncorrelated anomalies are captured in the first 4 principal components projection. The Q-statistics (δ^2) are also shown in these plots. From the lower plot, we found Q-statistic (δ^2) is sensitive to the detected correlated anomalies but not the uncorrelated anomalies. Based on these characteristics of correlated and uncorrelated anomalies in the normal subspace, and the drawbacks of Q-statistic, we apply the APD scheme to detect anomalies.

The link measurements \mathbf{y} 's projection onto normal subspace in Eqn. (4) can be written as:

$$\mathbf{SPE} \equiv \|\mathbf{y}_{no}\|^2 = \|\mathbf{P}\mathbf{P}^T\mathbf{y}\|^2, \mathbf{P} = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \dots, \mathbf{v}_k] \quad (5)$$

In APD [9], we use $\mu \pm s * \sigma$ as a threshold to define the norm of network health, where μ denotes the mean of measurements samples, σ is standard deviation of the measurements samples, and sensitivity s specifies the magnitude of the plateau change that may result when an anomaly event on a network path is to be triggered. Using this APD scheme threshold definition, we may consider the network measurements to be normal if,

$$\mu - s * \sigma \leq \mathbf{SPE} \leq \mu + s * \sigma \quad (6)$$

Now if we combine Eqn. (6) with Q-statistic, we formalize conditions for correlated and uncorrelated anomalies. We say correlated anomalies have occurred in the network measurements if,

$$\begin{cases} \mathbf{SPE} > \mu + s * \sigma & \text{and} & \mathbf{SPE} > \delta^2 \\ \delta^2 < \mathbf{SPE} < \mu - s * \sigma & & \end{cases} \quad (7)$$

And similarly, we conclude that uncorrelated anomalies have occurred in the network measurements if,

$$\begin{cases} \mathbf{SPE} < \mu + s * \sigma & \text{and} & \mathbf{SPE} < \delta^2 \\ \delta^2 > \mathbf{SPE} > \mu - s * \sigma & & \end{cases} \quad (8)$$

With the correlated and uncorrelated anomaly detection conditions formalized, we analyze the accuracy of our proposed anomaly detection scheme in Section VI.

V. DATA SANITY CHECKING

The efficacy of our proposed anomaly detection scheme relies heavily on the quality of the collected measurement samples. However, due to mis-calibration of measurement probes and potentially improper sampling (from the perspective of a monitoring objective such as rapid and accurate anomaly detection or accurate network weather forecasting [38]) in perfSONAR, the samples collected from multiple domains' measurement archives are not always worthy of analysis. In this section, we propose a two-pronged approach to sanitize perfSONAR measurement data: a reputation analysis scheme for collected samples, and a filter framework to intelligently prune the potentially misleading samples.

A. Reputation Analysis

In order to ascertain what features in a sample set of data qualify as 'good', we collected a considerable amount of perfSONAR one-way delay traces for different paths and different time periods. In any random collection that are publicly accessible, we observed some measurements exhibit non-periodic sampling pattern, i.e., these samples are either too dense or too sparse, and some are invalid due to faulty clock synchronization between measurement servers or data corruption (negative one-way delay values). Such improper sampling ultimately results in erroneous detections and consequently useless diagnosis/notifications (i.e., increased false alarms) using our proposed anomaly detection scheme [9]. Therefore, in the context of detecting and diagnosing potential correlated anomaly events within perfSONAR one-way delay traces, it is of paramount importance that the sample data has desired nature expected by the monitoring objective in terms of 2 aspects: *Sampling Pattern* and *Data Validity*.

To identify potentially misleading features of measured data, we propose a reputation-based data sanity checking scheme which analyzes the measurement samples for sampling pattern, and collected sample validity. This scheme involves reputation score evaluation for each measurement path, and "certainty" quantification of the entire measurement dataset through a propagation function as shown in Fig. 8. The certainty of detection strengthens the conclusions drawn about the nature and location of possible correlated anomaly events output by our anomaly detection scheme.

1) *Effect of sampling pattern*: Periodic, random, stratified random, and adaptive sampling are the most common sampling patterns in network performance measurements [38] to satisfy various monitoring objectives. Our APD algorithm is a real-time detector based on the time series' measurement data to detect correlated anomaly events. Thus, our algorithm requires the measurement data should be continuous and periodic in terms of the sampling time-intervals [39]. Recall from Section I that highly dense/sparse or irregular (i.e., long data collection gaps) can result in missed anomaly events and exponential anomaly detection time [9].

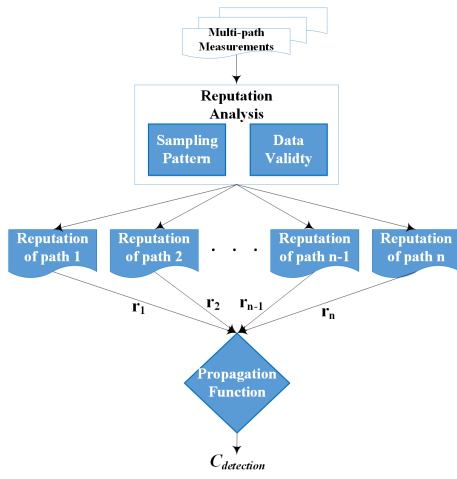


Fig. 8: Reputation-based scheme to evaluate the certainty of a correlated anomaly event detected by PCA-APD-Q-statistic

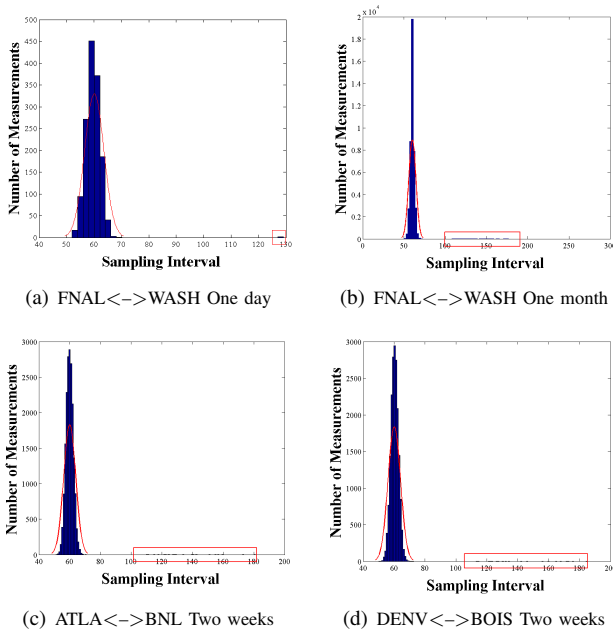


Fig. 9: Measurement sampling time-interval histograms for one-way delay perfSONAR traces

In order to conduct a deeper investigation on the sampling patterns of perfSONAR data, we collect perfSONAR one-way delay measurement data from different DOE lab and ESnet sites for different time periods. Figs. 9(a) and 9(b) show one such exemplar sampling time interval histogram for one-way delay measurements from DOE lab site FNAL to ESnet POP site WASH. From the figure, it is evident that the majority of sampling time-intervals are gathered in the one zone (marked by red curve) which suggest that the majority exhibits expected characteristic in terms of sampling pattern (i.e., good quality data) with outliers (marked by a red box) being abnormal. Similar characteristics were observed for other DOE lab and ESnet sites for different time periods as shown in Figs. 9(c) and 9(d). Such a pattern signifies that if a perfSONAR domain is in adequately calibrated, all the time intervals should exhibit the majority property.

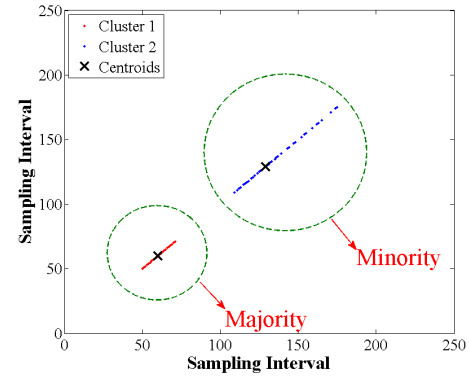


Fig. 10: K-Means Clustering to partition sampling time-intervals for separating major and minority clusters in measurement samples

In order to isolate the ‘good’ samples from the ‘not-so-good’ ones, we use the K-Means Clustering algorithm to partition the sampling time-interval majority and minority clusters. Fig. 10 shows the use of K-Means Clustering algorithm to partition one month duration of one-way delay measurements for FNAL \leftrightarrow WASH.

2) *Effect of data validity*: Measurement data in a few instances becomes invalid because of faulty clock synchronization and/or data corruption. For example, clock synchronization problem between measurement servers will cause the one-way delay values of OWAMP measurements to be negative; whereas data corruption will cause the value of delay to be ‘NaN’. As discussed earlier, incorporating such invalid data for anomaly detection analysis can lead to erroneous notifications. For a measurement data to be valid, the value of delay should be larger or equal than zero.

3) *Reputation and certainty quantification*: With sampling pattern and data validity being the two most important factors in deciding the quality of perfSONAR one-way delay data, we propose the reputation of any path i to be defined as:

$$r_i = 1 - \frac{N_i - n_i^{majority}}{N_i} - \frac{N_i - n_i^{valid}}{N_i} = \frac{N_i - (N_i - n_i^{majority}) - (N_i - n_i^{valid})}{N_i} \quad (9)$$

where N_i denotes the number of measurement samples in path i , n_i^{valid} denotes the number of valid data samples in path i , and $n_i^{majority}$ denotes the number of samples in the majority zone of path i .

As the reputation score is specific to one measurement path, we still require a mean to translate the reputations of each path ($r_i \forall i \in N$) into a reputation score of the entire measurement data set. This measurement reputation score quantifies the certainty of the detection, which inherently guides a network operator to assess the true severity of the detected anomaly event. Although the NIST guidelines on measurement uncertainty quantification are not applicable for our measurement data reputation analysis, we use the NIST guidelines for measurement uncertainty propagation [40] in order to quantify the certainty of detection ($C_{detection}$) from measurement path reputation scores. As the measurement of different source-destination pairs are uncorrelated, i.e., there

are M mutually exclusive measurement observations, the corresponding certainty of detection is given by:

$$C_{\text{detection}} = \sum_{i=1}^M r_i^2 \quad (10)$$

where M is the total number of measurement paths. The above equation ensures positive and negative certainty of detection; thus clearly distinguishing the High and Low certainty values. Also, $C_{\text{detection}}$ monotonically increases for increasing number of paths with higher reputation scores and vice versa.

B. Filter Framework

Unlike correlated anomaly events, uncorrelated events being manifestations of network related faults at an external domain are harder to localize in the absence of complete network topology information. In order to investigate root-cause locations of such uncorrelated anomaly events without topology information, we propose a novel ‘‘Filter Framework’’ to which we pass the multi-path time-series measurement data through a series of filters. A collection of temporal and spatial filters are applied on the time-series data depending on the relative certainty of uncorrelated anomaly detection and a recursive PCA-APD-Q-statistic analysis is applied on the filter output.

1) *Destination to source conversion for high certainty uncorrelated anomaly events:* When uncorrelated anomaly events are detected with high certainty, we apply temporal filters on the time-series data to find the paths with uncorrelated anomaly event timestamp. The temporal filters consists of two parts: *filtration* and *measurements transformation*. First, we transform all the time-series measurements into a matrix whose columns list the measurements information such as timestamp, measurement value, and detection results. Each measurement path is transformed into a matrix resulting multi-path measurements generating matrices. Next, linear search is performed on each matrix using the timestamp as the index to select the row having an anomaly event around the same time window. The output of the filters are the individual paths suspected to be the responsible uncorrelated event path.

Upon filtering the paths, a new set of measurement data are collected for the same time period as the original, with the destinations of each path now being the new sources. If correlated anomaly events are detected with high certainty upon analysis of the new samples using our scheme, we can localize the original uncorrelated anomaly events at the respective destination. Otherwise, the original uncorrelated anomaly events were caused due to some abnormal network behavior at one or many points along the paths other than the sources and destinations.

2) *Pruning misleading data for low certainty uncorrelated anomaly events:* If uncorrelated anomaly events are detected with low certainty, we argue that the outcome of the analysis can be dubious due to potentially untrustworthy (low reputation) samples and the very existence of uncorrelated anomaly events maybe in question. Thus, in such cases, we apply spatial filters to detect and intelligently prune potentially misleading samples, and re-analyze the new trimmed sample set for anomaly events.

The spatial filtration and the subsequent pruning are based on the reputation analysis of measurement data discussed in Section V-A. In order to ascertain what reputation is below-par and whether reputation is source-specific, we collect per-SONAR measurement archives from various DOE sites (e.g., FNAL, KANS, HOUS) for different time periods to analyze

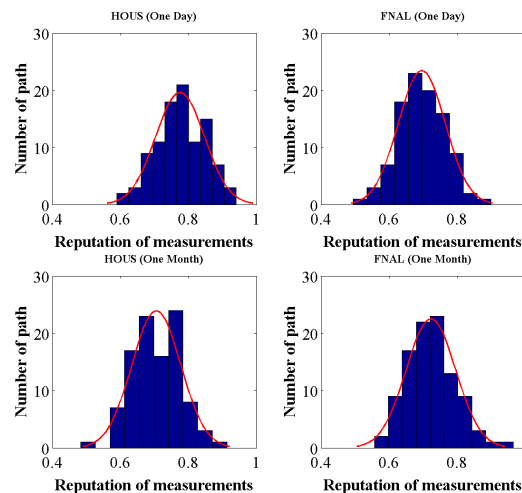


Fig. 11: One-way delay measurement sample reputation of HOUS and FNAL to other sites showing normal distribution

their historic distribution pattern. Jarque-Bera test [41] on their reputation distribution reveals that for every DOE source site, the measurements’ reputation unanimously follows normal distribution. In Fig. 11, we show two exemplar measurement data reputation histograms of ESnet sites HOUS and FNAL to other sites for different time periods following a normal distribution.

Pruning misleading data can be tricky as single-dimension reputation-centric pruning may lead to excessive thinning of measurement samples, i.e., we are left with too few samples to analyze effectively. Thus, we take a multi-dimensional approach of pruning based on both reputation and sample population. Depending on the size of the population we might want to keep at least $K\%$ of the entire sample population; if the population is big, value of K can be smaller and vice versa. As any sample population follows normal distribution ($N(\mu, \delta^2)$ with μ being the mean and δ being the standard deviation), we can calculate the baseline reputation x for keeping $K\%$ of the entire sample population with the relation $P(X \geq x) = K/100$ and Z score formula $Z = \frac{x-\mu}{\delta}$. Thus, any sample with reputation less than baseline x is pruned and the new sample set is re-analyzed for anomaly events. For example, in case of a relatively low number of total available samples for detection and diagnosis, the network operator may decide to keep at least 80% ($K = 80$) of the entire sample population, even if some samples are not good enough. The reasoning might be that for very small sample size, it is more prudent to prune a few worst samples rather than ending up with very few samples for actual detection. In such a case, the baseline reputation x is calculated from normal distribution curve with the relation $P(X \geq x) = 0.8$.

In another scenario with a very high number of samples, the network operator can be more ruthless about the quality of the collected samples and may decide to keep only the best 30% ($K = 30$) of the entire sample population for effective detection and diagnosis. In such a case, the baseline reputation x can be similarly calculated from normal distribution curve with the relation $P(X \geq x) = 0.3$.

3) *Implementation of our Proposed Scheme:* We demonstrate the implementation of the proposed scheme in our

Narada Metrics [13] framework for actual perfSONAR one-way delay measurements. Narada Metrics features perfSONAR extensions that can analyze network performance via monitoring-objective directed sampling, and generates performance trend reports and notifies anomaly events to communities subscribed to measurement archives in perfSONAR. Before this work, Narada Metrics used just the APD scheme for path-level detection of uncorrelated anomaly events amongst collected measurements over a specified user time range. We extend the Narada Metrics path-level analysis in this work to a network-wide level by adding PCA and temporal/spatial filters. Fig. 12 shows Narada Metrics performance plot of anomaly events with temporal correlation between two traces of one-way delay measurements. For each trace, we apply our APD scheme to detect anomaly events, which are shown with annotations for network operators to perform further drill-down analysis.

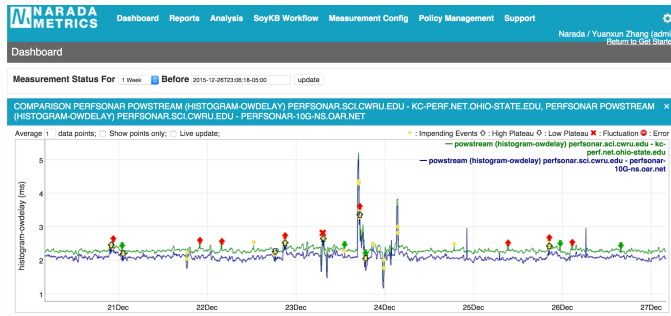


Fig. 12: Narada Metrics temporally correlated anomaly event notification

Fig. 13 shows how the paths responsible for uncorrelated anomaly events are filtered using the PCA anomaly event timestamp. As shown in “Measurements Data Description” of Fig. 13, we collect 5 perfSONAR one-way delay measurements traces and upon PCA-APD-Q-statistic analysis, the uncorrelated anomaly event timestamps are calculated as 2015-04-21T18:40:00+00:00 and 2015-04-21T15:20:00+00:00. The temporal filtration function uses these timestamps as inputs to filter the traces responsible for the anomaly events as the “Filtration Results”. We implement spatial filters in our Narada Metrics framework as shown in the Fig. 14. For this feature, we collected thousands of measurement traces as training dataset to generate the Normal distribution of measurement data reputations. As shown in the illustration, a network operator can prune below-par samples in terms of measurement paths by making sure that at least a desired percentage (e.g., 80%) of the entire dataset is retained. The algorithm computes the reputation baseline according to the Normal distribution as shown in “Reputation baseline”. The output of the filtration processes are the paths to be eliminated as shown in the “Filtration Results” window.

VI. EVALUATION AND RESULTS

A. Anomaly Detection and Isolation with Synthetic data

To demonstrate the PCA-APD-Q-statistic scheme’s accuracy of detection and isolation, we plot the anomaly detection performance of the existing PCA-APD [7], PCA-SPD [36], and PCA-Q-statistic [11, 12] (without APD) schemes with different datasets and different number of correlated and

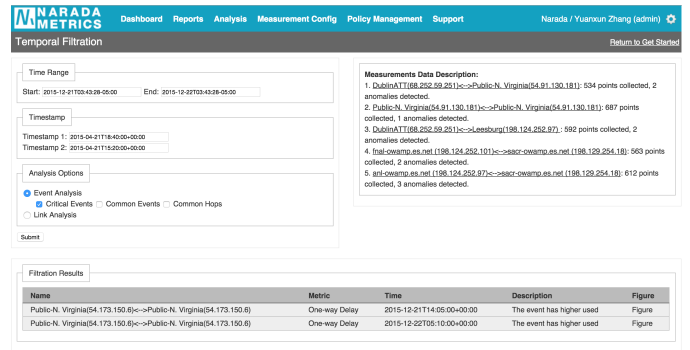


Fig. 13: Narada Metrics Temporal Filter implementation

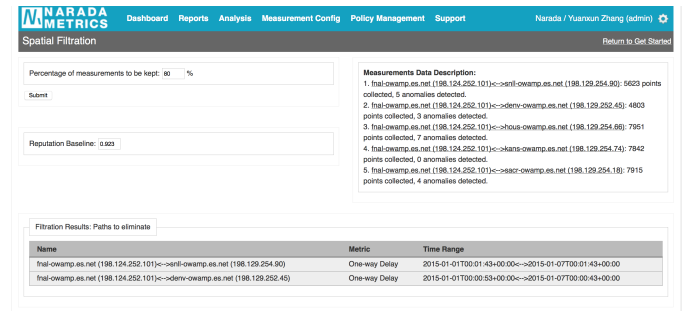


Fig. 14: Narada Metrics Spatial Filter implementation

uncorrelated anomaly events. Recall that the Q-statistic is a statistic test to detect threshold-crossing samples. To adapt the Q-statistic into a plateau detector, we look for 7 (same trigger count in APD and SPD) consecutive threshold-crossing to classify it as a plateau event.

1) *Evaluation methodology*: For this, we generate synthetic trace data described earlier in Section IV-B. We randomly generate the one week dataset, and inject different number of correlated and uncorrelated anomaly events into the dataset to compare the detection accuracy. The synthetic data is carefully generated to closely mimic the actual perfSONAR one-way delay measurement traces, as shown in the Fig. 15.

In order to inject correlated anomaly events, we first generate 6 traces and then inject anomaly events in those traces at the same time. We also inject events at random times as uncorrelated anomaly events. The percentage of anomaly events in each trace ($\rho_{anomaly}$) vary from 0.1%-1% of the total sample population for each trace. The magnitudes of anomaly events ($r_{magnitude}$) vary from 10% - 60% over normal measurements with higher magnitudes causing sharper spikes.

2) *Evaluation metrics*: We evaluate the anomaly event detection accuracy of our PCA-APD scheme and compare it with PCA-SPD and PCA-Q-statistic schemes using three well known detection evaluation metrics, viz., Accuracy, False Positive Rate (FPR), and False Negative Rate (FNR). These metrics are defined as follows:

$$Accuracy = \frac{\sum True Positives + \sum True Negatives}{\sum Anomaly Events}$$

$$FPR = \frac{\sum False Positives}{\sum Anomaly Events}$$

$$FNR = \frac{\sum False Negatives}{\sum Anomaly Events}$$

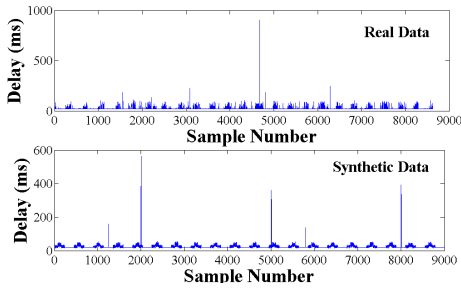


Fig. 15: Data sample comparison between real data and synthetic data

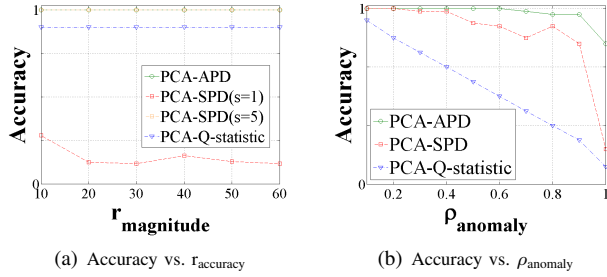


Fig. 16: Detection accuracy comparison

3) *Anomaly detection evaluation results:* In Figs. 16(a) and 16(b), we compare the accuracy of the schemes. In Fig. 16(a), we show the nature of detection accuracy with $r_{\text{magnitude}}$. We observe, APD and SPD with high sensitivity ($s=5$), exhibit 100% detection accuracy. However, a less sensitive SPD ($s=1$) suffers from poor accuracy due to erroneous detection of especially uncorrelated anomaly events that are difficult to detect. In Fig. 16(b), we compare the scheme's accuracy against ρ_{anomaly} . We observe that APD performs robustly against high-density of anomaly events in the samples. Both SPD and Q-statistic suffer from erroneous detection of uncorrelated anomaly events, whose number increases with higher ρ_{anomaly} .

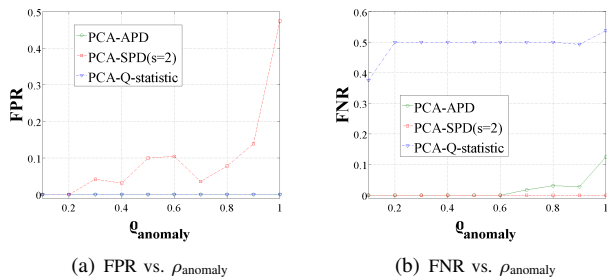


Fig. 17: False alarm rate comparison

In Fig. 17(a), we compare the scheme's performances in terms of FPR against ρ_{anomaly} . We observe that the APD and Q-statistic do not react to anomaly density, however the performance of SPD ($s=2$) rapidly deteriorates due to static threshold settings, thus causing more false alarms of anomaly events. In the Fig. 17(b), we can see that SPD performs the best for the default over-sensitivity of detection. The Q-statistic misses all the uncorrelated anomaly events and APD exhibits

few false alarms at a higher anomaly density.

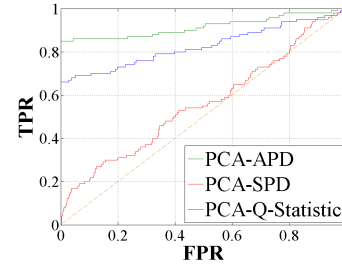


Fig. 18: Receiver operating characteristic comparison

A meta-perspective of the scheme's relative performances are shown in Fig. 18 through a Receiver Operating Characteristic (ROC) plot. The ROC plot indicates the True Positive Rate (TPR) performance of a scheme against FPR with performance curve above the 45° denoting better performance. We see that our PCA-APD scheme clearly outperforms the other two schemes with a higher accuracy rate than false alarm rate indicating that the PCA-APD is the best among the contending three schemes of anomaly detection.

4) *Anomaly isolation evaluation results:* Next, we analyze the performance of our proposed PCA-APD-Q-statistic in isolating correlated and uncorrelated anomaly events. In the Section IV-B, we explained how Q-statistic is a threshold to discern the correlated and uncorrelated anomaly events, where the anomaly events above the threshold are correlated anomaly events, and below are uncorrelated. However, we observed that the size of the measurement dataset plays an important role in the relative performance of Q-statistic.

In Fig. 19(a), we show how the Q-statistic successfully isolates correlated and uncorrelated anomaly events in a one week synthetic data set with 5 correlated and 2 uncorrelated anomaly events injected. However, the same 7 anomalies for a two weeks long dataset are not successfully distinguished with both the uncorrelated anomaly events being detected wrongly as correlated, as shown in the Fig. 19(b). The performance improves again upon injecting more correlated anomaly events in the dataset, with 100% accuracy reached when 3 more correlated anomaly events are injected (i.e., total 8 correlated and 2 uncorrelated), as shown in the Fig. 19(c). Similar characteristics are observed in month long data and higher. The reason behind such characteristics is the dynamic trigger demotion of Q-statistic threshold by predominant uncorrelated anomaly events in the dataset since the magnitude of the uncorrelated anomaly events are much smaller in the data projection. With more correlated anomaly events in the dataset, such effects are negated and the Q-statistic performs better.

In order to gauge the accuracy of isolation in an unknown data set, the proximity of a correlated anomaly event in higher principal components to the Q-statistic threshold may prove to be more useful. To this end, we conducted rigorous experiments with synthetic datasets of different sizes, and concluded that any correlated anomaly detected within the range of $0 - 2 \times 10^5$ from the Q-statistic threshold in second principal component data projection, called the 'Grey zone' is inconclusive and may need further investigation with artificially injected correlated anomaly events in the dataset.

5) *Detection and isolation evaluation summary:* The detection accuracy evaluation showcases the improved accuracy

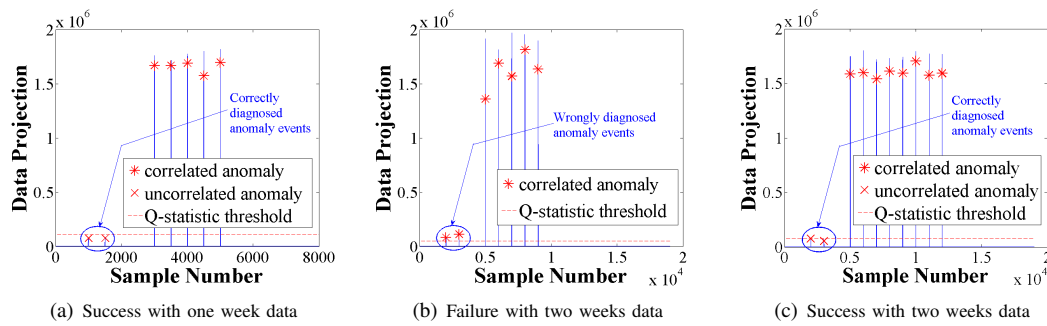


Fig. 19: Accuracy of Q-statistic to successfully isolate correlated and uncorrelated anomaly events for different dataset sizes

of using our proposed PCA-APD-Q-statistics anomaly event detection technique over existing schemes in the absence of complete network topology information. Moreover, our proposed scheme also result in lower false alarms rates over existing schemes for different densities of anomaly events in the measurement traces. In anomaly isolation evaluation, we presented only a small subset of total number of experiments performed with synthetic data. The results overwhelmingly demonstrate how our proposed scheme successfully isolates correlated events from uncorrelated events. The results also help us generate bounds on the relative ratio of correlated and uncorrelated anomaly events for pertinent isolation.

B. Case Studies with actual perfSONAR Data

In this section, we validate the use of our proposed anomaly detection and certainty diagnosis scheme to analyze correlated and uncorrelated anomaly events at the network-wide level using source-site information within actual perfSONAR traces. The datasets in the following case studies consist of plateau anomalies such as persistent increase and other anomaly events such as intermittent bursts and dips. We consciously ignore intermittent burst and dip events because these types of anomalies are generally caused by user behavior, and are not of interest to network operators for routine monitoring and bottleneck troubleshooting. All of the actual perfSONAR traces correspond to one-way delay measurements collected between DOE lab sites such as FNAL (Fermi National Accelerator Laboratory), SLAC (SLAC National Accelerator Laboratory), ORNL (Oak Ridge National Laboratory), and ESnet 100G hubs ATLA (Atlanta), STAR (StarLight), and SUNN (Sunnyvale).

We perform a select set of case studies using both short-term and long-term traces to demonstrate different functionalities of our proposed scheme, such as: (i) anomaly detection, (ii) anomaly correlation identification, (iii) anomaly detection certainty evaluation, (iv) diagnosing potential uncorrelated anomaly location through destination to source conversion, and (v) effect of pruning misleading samples to increase detection certainty. The purpose of collecting samples for different time periods and time lengths below is to demonstrate the effectiveness of our scheme for both short-term and long-term measurement objectives using different sample size populations.

1) Case Study I: Location isolation with one month data:

As discussed in Section IV, the PCA-APD-Q-statistic scheme can accurately detect correlated and uncorrelated anomaly events with low false alarm rates, and PCA-with-Q-statistic scheme can completely accurately detect correlated anomaly

events (however, it misses all the uncorrelated anomaly events), we leverage the Q-statistic within our PCA-APD scheme as a sure way to accurately identify all the correlated anomaly events. To illustrate with further evidence, we used the actual perfSONAR one-month long traces from ESnet site ATLA to 6 other DOE lab sites as shown in Table I. Fig. 20(a) shows one anomaly that is detected in the actual perfSONAR traces by the PCA-APD-Q-statistic scheme using the first principal component. If we assume all correlated anomalies are captured in the first principal component, and the uncorrelated anomalies are captured in the rest of principal components, we may mis-identify correlated anomaly events in certain situations. Consequently, as shown in Fig. 20(b), our scheme identifies the anomaly above the Q-statistic as a correlated anomaly event with a high certainty of detection 0.9979 showing high confidence on the collected samples. From the above analysis, we can conclude a correlated anomaly occurred in a local domain (i.e., within ATLA) at 22:29:11-22:38:34 time period. In order to validate this detection, we checked each of the traces using just the APD scheme to detect anomaly events in each trace. We found six traces to have the anomaly events at the same time windows.

The Case Study I results show how a network operator can use the PCA-APD-Q-statistic scheme to detect and analyze the correlation among anomaly events without having complete topology information.

TABLE I: Case Study I perfSONAR traces description

Source ↔ Destination	Time Range (Start - End)
atla-owamp.es.net ↔ boise-owamp.es.net	2014-10-01 00:00:55 - 2014-10-30 23:59:43
atla-owamp.es.net ↔ chic-owamp.es.net	2014-10-01 00:00:30 - 2014-10-30 23:59:04
atla-owamp.es.net ↔ fnal-owamp.es.net	2014-10-01 00:01:00 - 2014-10-30 23:59:52
atla-owamp.es.net ↔ hous-owamp.es.net	2014-10-01 00:00:48 - 2014-10-30 23:59:52
atla-owamp.es.net ↔ nersc-owamp.es.net	2014-10-01 00:00:44 - 2014-10-30 23:59:33
atla-owamp.es.net ↔ wash-owamp.es.net	2014-10-01 00:00:16 - 2014-10-30 23:59:38

2) Case Study II: Certainty diagnosis with one month data:

In order to ensure that the accuracy of location isolation using the proposed scheme also holds for other paths, we collect one month measurements from ESnet site STAR to 8 DOE lab sites as shown in Table II. Upon PCA-APD-Q-statistic analysis and data sanity checking, we detected 3 correlated anomaly events (as shown in the Fig. 19(c)), however with low certainty of detection, 0.4669. As we already established that source-side faults lead to correlation, local domain (i.e., STAR) is most

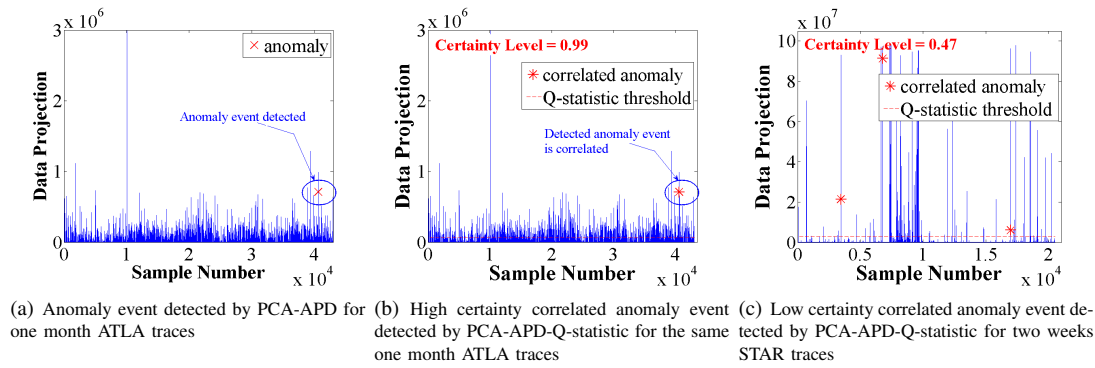


Fig. 20: Case Study I and II: Location Isolation and Certainty Diagnosis

likely to be responsible for the anomaly events, albeit with low detection certainty.

Thus, Case Study II results show how a network operator can handle scenarios where the measurement samples with misleading features can lead to low certainty of detection. For e.g., given the Case Study II certainty of detection being low, the network operator can prioritize fixing Case Study I issues over Case Study II issues when troubleshooting bottlenecks.

TABLE II: Case Study II perfSONAR traces description

Source ↔ Destination	Time Range (Start - End)
star-owamp.es.net ↔ bois-owamp.es.net	2014-10-01 00:00:55 - 2014-10-30 23:56:35
star-owamp.es.net ↔ elpa-owamp.es.net	2014-10-01 00:03:00 - 2014-10-30 23:59:36
star-owamp.es.net ↔ ga-owamp.es.net	2014-10-01 00:00:56 - 2014-10-30 23:59:31
star-owamp.es.net ↔ kans-owamp.es.net	2014-10-01 00:03:00 - 2014-10-30 23:58:36
star-owamp.es.net ↔ llnl-owamp.es.net	2014-10-01 00:00:20 - 2014-10-30 23:59:47
star-owamp.es.net ↔ sdsc-owamp.es.net	2014-10-01 00:05:00 - 2014-10-30 23:56:35
star-owamp.es.net ↔ sunn-owamp.es.net	2014-10-01 00:00:25 - 2014-10-30 23:55:35
star-owamp.es.net ↔ wash-owamp.es.net	2014-10-01 00:00:57 - 2014-10-30 23:58:36

3) *Case Study III: Root cause diagnosis of low certainty uncorrelated anomaly event with one month data:* We already discussed that detected uncorrelated anomaly events in most cases are a result of faults outside the local domain. However, in a measurement setup with multiple paths with high topological correlation (i.e., common hops), such uncorrelated anomaly events on individual paths at different timestamps can be isolated to be manifestations of faults at a destination domain. In the absence of network topology information, destination-to-source transformation of the path having uncorrelated anomaly event and subsequent measurement data collection for the same time period can lead to accurate location diagnosis. To test this rationale, we collect month-long measurements from DOE lab site FNAL to 16 other DOE lab and ESnet sites (as shown Table III) with high topological correlation among the paths. When we applied our PCA-APD-Q-statistic anomaly detection scheme and sanity checking, we detected an uncorrelated anomaly event with high certainty as shown in the Fig. 21(a). Upon applying filters, we found that the anomaly occurred at around 2014-10-29 05:46:53 on the path FNAL to SLAC in Table III.

Thus, with the expectation to find a correlated anomaly event with high certainty around the same time, we collected perfSONAR data with SLAC as source to 6 other DOE sites as shown in Table IV. Upon analysis, we detected a correlated anomaly event with high certainty as shown in the Fig. 21(b) occurring at around 2014-10-29 05:40:24, thus validating our original claim of detecting correlation. The Case Study III results help a network operator to successfully use our proposed temporal filter (for destination to source conversion) in diagnosing and isolating uncorrelated anomaly events during root-cause analysis, which otherwise would be challenging in practice without complete topology information.

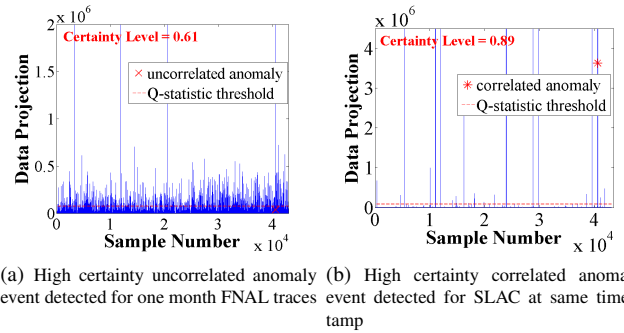


Fig. 21: Case Study III: Destination to Source Conversion

4) *Case Study IV: Pruning misleading measurement samples with one week data:* To validate the claim of pruning misleading data can improve detection accuracy, we collected short term (one week) perfSONAR samples from ORNL to 6 other DOE sites as shown in Table V. The PCA-APD analysis fails to find any anomaly events as shown in the Fig. 22(a). However, the certainty of detection was 0.83 which was less than the historical average of ORNL specific data.

To avoid excessive thinning of measurement samples before pruning the ill-reputed samples, we analyzed the historical ORNL measurements data and estimated with Z score formula (discussed in Section V-B2) to prune samples below 0.8798 (minimum reputation of 80% sample population). Upon filtering, we eliminated the trace from ORNL to SNLA (shown in red) with reputation 0.07 that was impacting the detection outcome and executed recursive PCA-APD-Q-statistic on the new sample set. Upon analysis, we found an uncorrelated anomaly event (as shown in Fig. 22(b)), undetected previously.

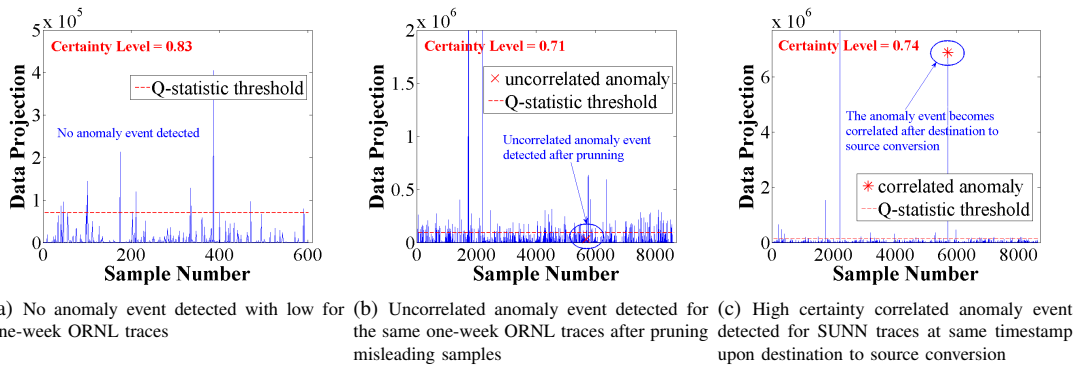


Fig. 22: Case Study IV: Pruning Misleading Measurement Samples

TABLE III: Case Study III perfSONAR traces description with FNAL as source

Source ↔ Destination	Time Range (Start - End)
fnal-owamp.es.net ↔ bois-owamp.es.net	2014-10-01 00:00:31 - 2014-10-30 23:59:47
fnal-owamp.es.net ↔ ga-owamp.es.net	2014-10-01 00:00:08 - 2014-10-30 23:59:45
fnal-owamp.es.net ↔ hous-owamp.es.net	2014-10-01 00:00:04 - 2014-10-30 23:59:58
fnal-owamp.es.net ↔ lasv-owamp.es.net	2014-10-01 00:00:50 - 2014-10-30 23:59:33
fnal-owamp.es.net ↔ lbl-owamp.es.net	2014-10-01 00:00:29 - 2014-10-30 23:59:00
fnal-owamp.es.net ↔ nersc-owamp.es.net	2014-10-01 00:00:25 - 2014-10-30 23:59:00
fnal-owamp.es.net ↔ newy-owamp.es.net	2014-10-01 00:00:13 - 2014-10-30 23:59:58
fnal-owamp.es.net ↔ ornl-owamp.es.net	2014-10-01 00:00:36 - 2014-10-30 23:59:22
fnal-owamp.es.net ↔ pnwg-owamp.es.net	2014-10-01 00:00:32 - 2014-10-30 23:59:52
fnal-owamp.es.net ↔ pppl-owamp.es.net	2014-10-01 00:00:01 - 2014-10-30 23:59:40
fnal-owamp.es.net ↔ sacr-owamp.es.net	2014-10-01 00:00:01 - 2014-10-30 23:59:31
fnal-owamp.es.net ↔ sdsc-owamp.es.net	2014-10-01 00:00:21 - 2014-10-30 23:59:37
fnal-owamp.es.net ↔ slac-owamp.es.net	2014-10-01 00:00:57 - 2014-10-30 23:59:51
fnal-owamp.es.net ↔ snll-owamp.es.net	2014-10-01 00:00:12 - 2014-10-30 23:59:05
fnal-owamp.es.net ↔ sunn-owamp.es.net	2014-10-01 00:00:05 - 2014-10-30 23:59:34
fnal-owamp.es.net ↔ wash-owamp.es.net	2014-10-01 00:00:08 - 2014-10-30 23:59:49

TABLE IV: Case Study III perfSONAR traces description with SLAC as source

Source ↔ Destination	Time Range (Start - End)
slac-owamp.es.net ↔ bois-owamp.es.net	2014-10-01 00:00:45 - 2014-10-30 23:59:35
slac-owamp.es.net ↔ elpa-owamp.es.net	2014-10-01 00:00:47 - 2014-10-30 23:59:51
slac-owamp.es.net ↔ kans-owamp.es.net	2014-10-01 00:00:07 - 2014-10-30 23:59:55
slac-owamp.es.net ↔ pnwg-owamp.es.net	2014-10-01 00:00:40 - 2014-10-30 23:59:37
slac-owamp.es.net ↔ star-owamp.es.net	2014-10-01 00:00:17 - 2014-10-30 23:59:57
slac-owamp.es.net ↔ wash-owamp.es.net	2014-10-01 00:00:32 - 2014-10-30 23:59:58

We apply similar destination-to-source conversion to that of Case Study III to collect new measurement data (as shown in Table VI) for the path responsible for the uncorrelated anomaly event (SUNN). As expected, our analysis proved the existence of a correlated anomaly event with high certainty (as

shown in Fig. 22(c)) around the same time period. Through Case Study IV, we established that pruning misleading data using our proposed spatial filter can help a network operator to obtain higher detection certainty for further analysis on new trimmed sample sets. This in turn can reveal to new and more interesting features corresponding to anomaly events.

TABLE V: Case Study IV perfSONAR traces description with ORNL as source

Source ↔ Destination	Time Range (Start - End)
ornl-owamp.es.net ↔ slac-owamp.es.net	2015-01-01 00:00:01 - 2015-01-07 00:00:00
ornl-owamp.es.net ↔ snla-owamp.es.net	2015-01-01 00:00:14 - 2015-01-06 23:59:59
ornl-owamp.es.net ↔ srs-owamp.es.net	2015-01-01 00:00:31 - 2015-01-06 23:59:41
ornl-owamp.es.net ↔ star-owamp.es.net	2015-01-01 00:00:00 - 2015-01-06 23:59:18
ornl-owamp.es.net ↔ sunn-owamp.es.net	2015-01-01 00:00:50 - 2015-01-06 23:59:31
ornl-owamp.es.net ↔ wash-owamp.es.net	2015-01-01 00:00:04 - 2015-01-06 23:59:33

TABLE VI: Case Study IV perfSONAR traces description with SUNN as source

Source ↔ Destination	Time Range (Start - End)
sunn-owamp.es.net ↔ bois-owamp.es.net	2015-01-01 00:00:09 - 2015-01-06 23:59:01
sunn-owamp.es.net ↔ elpa-owamp.es.net	2015-01-01 00:00:34 - 2015-01-06 23:59:53
sunn-owamp.es.net ↔ kans-owamp.es.net	2015-01-01 00:00:03 - 2015-01-06 23:59:04
sunn-owamp.es.net ↔ pnwg-owamp.es.net	2015-01-01 00:00:49 - 2015-01-06 23:59:33
sunn-owamp.es.net ↔ sdsc-owamp.es.net	2015-01-01 00:00:51 - 2015-01-06 23:59:15
sunn-owamp.es.net ↔ wash-owamp.es.net	2015-01-01 00:00:37 - 2015-01-06 23:59:42

VII. CONCLUSION


In this paper, we presented a novel PCA-based network-wide correlated anomaly detection scheme that: (i) uses principal component analysis to capture the maximum variance in a given multiple path measurement time-series, (ii) applies adaptive plateau detector (APD) to detect anomaly events with fused data transformation by PCA, (iii) leverages Q-statistic event correlation analysis in order to accurately filter out correlated and uncorrelated anomaly events, and (iv) quantifies certainty of such detection using an adaptive reputation-based data sanity checking that accounts for factors such as sampling pattern, sampling frequency, and sample validity.

With the strength of our prior work in developing APD's accurate uncorrelated anomaly detection algorithm, our proposed PCA-APD-Q-statistic scheme in this paper has the unique ability to detect both correlated and uncorrelated anomalies with high accuracy and low false alarms, in a timely manner. With event correlation analysis, our scheme is suitable for source-side anomaly localization to help network operators to diagnose the root-cause of bottlenecks, even when network topology information is not completely available. The proposed scheme is able to filter our potential misleading data to associate a level of certainty for each such detection and diagnosis claims.

We implemented our novel scheme in the form of perfSONAR extensions and performed extensive validation experiments with both synthetic trace data and actual perfSONAR trace data collected from DOE lab and ESnet hub sites. Specifically, we presented four case studies that validate the utility of our PCA-APD-Q-statistic and data sanity checking schemes. Our work in this paper can help network operators using perfSONAR dashboard, and scientists of data-intensive applications to isolate and diagnose bottlenecks with a degree of certainty. Further, it can foster effective troubleshooting in the context of root-cause analysis of correlated network-wide anomaly events from a meta-perspective.

REFERENCES

- [1] A. Hanemann, J. Boote, E. Boyd, J. Durand, L. Kudirimoti, R. Lapacz, M. Swany, S. Trocha, J. Zurawski, "PerfSONAR: A Service Oriented Architecture for Multi-Domain Network Monitoring", *Proc. of Service Oriented Computing*, 2005.
- [2] J. B. Bottum, R. Marinshaw, H. Neeman, J. Pepin, J. B. von Oehsen, "The Condo-of-Condos", *Proc. of XSEDE*, 2013.
- [3] P. Kanuparth, D. Lee, W. Matthews, C. Dovrolis, S. Zarifzadeh, "Pythia: Detection, Localization, and Diagnosis of Performance Problems", *IEEE Communications Magazine*, vol. 51, no. 11, pp. 55-62, 2013.
- [4] C. Logg, L. Cottrell, "Experiences in Traceroute and Available Bandwidth Change Analysis", *Proc. of ACM SIGCOMM Network Troubleshooting Workshop*, 2004.
- [5] P. Callyam, M. Swany, "Research Challenges in Future Multi-domain Network Performance Measurement and Monitoring", *ACM Computer Communications Review* 45, pp. 29-34, 2015.
- [6] A. Mahimkar, J. Yates, Y. Zhang, A. Shaikh, J. Wang, Z. Ge, C. Ee., "Troubleshooting Chronic Conditions in large IP Networks", *Proc. of ACM CoNEXT*, 2008.
- [7] P. Callyam, M. Dhanapalan, M. Sridharan, A. Krishnamurthy, R. Ramnath, "Topology-Aware Correlated Network Anomaly Event Detection and Diagnosis", *Springer J. of Network and Systems Management*, Vol., 22, pp. 208-234, 2013.
- [8] V. Paxson, "Strategies for sound internet measurement", *Proc. of ACM IMC*, 2004.
- [9] P. Callyam, J. Pu, W. Mandrawa, A. Krishnamurthy, "OnTimeDetect: Dynamic Network Anomaly Notification in perfSONAR Deployments", *Proc. of IEEE MASCOTS*, 2010.
- [10] J. Jackson, G. Mudholkar, "Control Procedures for Residuals Associated with Principal Component Analysis", *Technometrics*, 1979.
- [11] A. Lakhina, M. Crovella, C. Diot, "Diagnosing Network-Wide Traffic Anomalies", *Proc. of ACM SIGCOMM*, 2004.
- [12] A. Soule, K. Salamian, N. Taft, "Combining Filtering and Statistical Methods for Anomaly Detection", *Proc. of ACM IMC*, 2005.
- [13] Narada Metrics: Software-defined measurement and monitoring. [Online - Last Accessed 01/07/2016]. <https://www.naradametrics.com>
- [14] R. Mahajan, N. Spring, D. Wetherall, T. Anderson, "User-level Internet Path Diagnosis", *Proc. of ACM SOSP*, 2003.
- [15] M. Zhang, C. Zhang, V. Pai, L. Peterson, R. Wang, "PlanetSeer: Internet Path Failure Monitoring and Characterization in Wide Area Services", *Proc. of ACM USENIX NSDI*, 2004.
- [16] H.V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, A. Venkataramani, "iPlane: An Information Plane for Distributed Services", *Proc. of ACM USENIX OSDI*, 2006.
- [17] D. R. Choffnes, F. E. Bustamante, Z. Ge, "Crowdsourcing Service-level Network Event Monitoring", *Proc. ACM SIGCOMM*, 2010.
- [18] F. Simmross-Wattenberg, J. Asensio-Perez, P. Casaseca-de-la-Higuera, M. Martin-Fernandez, I. A. Dimitriadis, C. Alberola-Lopez, "Anomaly Detection in Network Traffic Based on Statistical Inference and alpha-Stable Modeling", *IEEE Trans. on Dependable and Secure Computing*, Vol. 8, No. 4, pp. 494-509, 2011.
- [19] Y. Zhu, B. Helsen, J. Rexford, A. Siganporia, S. Srinivasan, "LatLong: Diagnosing Wide-Area Latency Changes for CDNs", *IEEE Trans. on Network and Service Management*, Vol. 9, No. 3, pp. 333-345, 2012.
- [20] M. Garbaoui, F. Paolucci, A. Giorgetti, B. Martini, P. Castoldi, "Effective Statistical Detection of Smart Confidentiality Attacks in Multi-Domain Networks", *IEEE Trans. on Network and Service Management*, Vol. 10, No. 4, pp. 383-397, 2013.
- [21] S. Gillani, M. Demirci, E. Al-Shaer, M. H. Ammar, "Problem Localization and Quantification Using Formal Evidential Reasoning for Virtual Networks", *IEEE Trans. on Network and Service Management*, Vol. 11, No. 3, pp. 307-320, 2014.
- [22] Y. Zhou, G. Hu, "Network-wide Anomaly Detection Based on Router Connection Relationships", *Proc. of IEICE Transactions*, 2011.
- [23] P. Barford, N. Duffield, A. Ron, J. Sommers, "Network Performance Anomaly Detection and Localization", *Proc. IEEE INFOCOM*, 2009.
- [24] H. Yan, A. Flavel, Z. Ge, A. Gerber, D. Massey, C. Papadopoulos, H. Shah, J. Yates, "Argus: End-to-end Service Anomaly Detection and Localization from an ISP point of view", *Proc. IEEE INFOCOM*, 2012.
- [25] A. Abdelkefi, Y. Jiang, B.E. Helvik, G. Bizcok, A. Calu, "Assessing the Service Quality of an Internet Path through End-to-end Measurement", *Elsevier Computer Networks*, Vol. 70, pp. 30-44, 2014.
- [26] A. Botta, A. Pescape, G. Ventre, "Quality of Service Statistics over Heterogeneous Networks: Analysis and Applications", *European Journal of Operational Research*, Vol. 191, No. 3, pp. 1075-1088, 2008.
- [27] P. Marchetta, P. Merindol, B. Donnet, A. Pescape, J. Pansiot, "Topology Discovery at the Router Level: A New Hybrid Tool Targeting ISP Networks", *IEEE JSAC*, Vol. 29, No. 9, pp. 1776-1787, 2011.
- [28] L. Zonglin, H. Guangmin, Y. Xingmiao, Y. Dan, "Detecting Distributed Network Traffic Anomaly with Network-Wide Correlation Analysis", *Proc. of EURASIP J. Adv. Signal Process*, Vol. 2, pp. 1-11, 2009.
- [29] Q. Ding, E. D. Kolaczyk, "A Compressed PCA Subspace Method for Anomaly Detection in High-Dimensional Data", *IEEE Trans. on Information Theory*, Vol. 59, No. 11, pp. 7419-7433, 2013.
- [30] Y. Zhang, P. Callyam, S. Debroy, M. Sridharan, "PCA-based Network-wide Correlated Anomaly Event Detection and Diagnosis", *Proc. of IEEE DRCN*, 2015.
- [31] M. Marvasti, A. Poghosyan, A. Harutyunyan, N. Grigoryan, "An Enterprise Dynamic Thresholding System", *Proc. of ICAC*, 2014.
- [32] Y. Tang, E. Al-Shaer, K. Joshi, "Reasoning under Uncertainty for Overlay Fault Diagnosis", *IEEE Trans. on Network and Service Management*, Vol. 9, No. 1, pp. 34-47, 2012.
- [33] G. Zacharia, "Trust Management through Reputation Mechanisms", *Proc. of Workshop in Deception, Fraud, and Trust in Agent Societies*, 1999.
- [34] Y. Wang, J. Vassileva, "Trust and Reputation Model in Peer-to-Peer Networks", *Proc. of IEEE P2P*, 2003.
- [35] D. Chadwick, S. Otenko, W. Xu, "Adding Distributed Trust Management to Shibboleth", *Proc. of NIST Annual PKI Workshop*, 2005.
- [36] A. McGregor, H-W. Braoun, "Automated Event Detection for Active Measurement Systems", *Proc. of PAM*, 2001.
- [37] P. Brockwell, R. Davis, "Introduction to Time Series and Forecasting", 2nd edition, Springer ISBN 978-0-387-21657-7, 2012.
- [38] P. Callyam, L. Kumarasamy, C. -G.Lee, F. Ozguner, "Ontology-based Semantic Priority Scheduling for Multi-domain Active Measurements", *Springer J. of Network and Systems Management (JNSM)*, Vol. 22, pp. 331-365, 2014.
- [39] D. Tammaro, S. Valenti, D. Rossi, A. Pescape, "Exploiting packet-sampling measurements for traffic characterization and classification", *ACM IJNM*, Vol. 22, No. 6, pp. 452-476, 2012.
- [40] B. N. Taylor, C. E. Kuyatt, "Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results", *NIST Technical Note 1297*, 1994.
- [41] C. M. Jarque, A. K. Bera, "A Test for Normality of Observations and Regression Residuals", *International Statistical Review*, Vol. 55, No. 2, pp. 163-172, 1987.



Yuanxun Zhang received his BE degree from Southwest Jiaotong University, China, in 2006. He is currently pursuing his PhD degree in University of Missouri-Columbia. His research interests include network performance monitoring, software-defined networking, and big data analytics.



Saptarshi Debroy received his PhD degree in Computer Engineering from University of Central Florida in 2014, MTEch degree from Jadavpur University, India in 2008, and BTech degree from West Bengal University of Technology, India in 2006. He is currently a Post Doctoral Fellow at University of Missouri-Columbia. His current research interests include cloud computing, Big Data networking, and cognitive radio networks.



Prasad Callyam received his MS and PhD degrees from the Department of Electrical and Computer Engineering at The Ohio State University in 2002 and 2007, respectively. He is currently an Assistant Professor in the Department of Computer Science at University of Missouri-Columbia. His current research interests include distributed and cloud computing, computer networking, and cyber security. He is a Senior Member of IEEE.