

Network Measurement Recommendations for Performance Bottleneck Correlation Analysis

Yuanxun Zhang, Saptarshi Debroy, Prasad Calyam

University of Missouri-Columbia

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

Abstract—Multi-domain network performance monitoring (NPM) federations, such as perfSONAR rely on collaborative measurement intelligence to identify network anomaly events and diagnose performance bottlenecks affecting data-intensive science applications. In this paper, we present a novel measurement recommendation scheme to assist network operators and application users by recommending pertinent samples from a pool of measurement data involving multiple domains to detect and troubleshoot correlated network anomaly events. The recommendations are based on the principles of content-based filtering. Such recommendations are complimented with Bayesian Inference based domain reputation meta-information to strengthen the veracity information of the recommended traces. Using actual long-term and short-term perfSONAR traces, we analyze recommendation results and show: a) how the content-based filter recommends the most pertinent traces based on their attributes, and b) the time-variant characteristics of domain reputation. Finally, using synthetic traces, we show the effectiveness of our proposed measurements recommendation scheme in accurately identifying anomaly events for an exemplar use case, and also show how our content filter based recommendation scheme performs better in terms of false alarms in comparison to: a) recommendations that consider partial trace features for filtering, and b) greedy recommendation approaches based on random trace selection.

Index Terms—Network Performance Monitoring, perfSONAR, Measurement Recommendations, Content-based Filtering.

I. INTRODUCTION

Scientific communities that support large-scale data transfers have extensively deployed multi-domain Network Performance Monitoring (NPM) federations that use passive and active measurements for monitoring and troubleshooting network bottlenecks. Among these NPM federations, perfSONAR [1] is the most widely instrumented framework within academia and industry (more than 1400 instances worldwide). It uses tools, such as Ping, Traceroute, OWAMP (for one-way delay measurements), and BWCTL (for TCP/UDP throughput measurements) to collect measurements such as end-to-end delay, jitter, loss, and bandwidth. With growing perfSONAR deployments within multi-domain federations, the initial focus of intra-campus network monitoring has shifted towards end-to-end performance monitoring and troubleshooting of Big Data applications. In multi-site Big Data collaborations, the application traffic generated within a network (domain) traverses several other domains or autonomous systems (ASes) before reaching a destination. As a result, end-to-end performance

monitoring and troubleshooting becomes significantly harder for standalone perfSONAR measurement instances (Measurement Point Appliances or MPAs) installed at the source and destination domains.

Thus, for accurate detection of network anomaly events that impact end-to-end application performance and corresponding root cause diagnosis, measurement data sharing and collaboration across domains becomes essential for network operators to satisfy measurement analysis objectives. Such a collaboration is manifested through sharing of perfSONAR measurement data archives within data-intensive scientific communities such as, Large Hadron Collider (LHC), Energy Science Network (ESnet), and Internet2 for public access. Network operators and domain application users subscribe to the perfSONAR data from such communities and use relevant measurement traces from their archives to troubleshoot similar network problems. However, the current perfSONAR framework does not provide the network operators and application users necessary tools or services to filter out the most relevant data from vast archives of measurement data that can help them accurately detect and diagnose network events of interest [2]. Thus, operators in most cases depend on greedy analysis based on intuition-led random measurement trace selection (trial-and-error) that could result in erroneous detection and ineffective diagnosis.

Furthermore, using shared measurement data from different communities for network anomaly event detection may not always prove useful as the measurement samples collected from such community perfSONAR deployments could have measurement mis-calibration issues, or issues such as invalid (e.g., negative one-way delay values due to faulty clock synchronization) or missing data. Such measurement data “veracity” issues result in erroneous features [3], or too dense/sparse or irregular (i.e., long data collection gaps) measurement sampling frequency. This in turn can lead to missed events or even exponential anomaly detection time [4]. In our earlier work [5], we have shown how solving the measurement data “veracity” problem can profoundly impact the accuracy of anomaly event detection analysis. However, with the wide adoption of perfSONAR measurement subscriptions within communities, the instantaneous quality of the measurement traces needs to be complimented with communities’ or domains’ long-term “confidence indicators” in terms of sanity of the measurement data being generated. Such “confidence indicators” are essential for operators and application users in order to subscribe to the communities that are *reputed* for good measurement practices. This can consequently lead to use of sanitized measurement data, which facilitates accurate detection and subsequent troubleshooting upon analysis.

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.

In this paper, we propose a “content filter” based measurement recommendation scheme that recommends pertinent measurement traces from a pool of candidate samples to assist network operators and application users. The proposed scheme ranks and recommends the most pertinent traces based on similarity matching with a target trace/path for which the operator/user needs help to perform some specific measurement correlation analysis. The similarity matching scheme: a) applies the principles of a content-based filtering technique by taking into account the candidate traces’ features (spatial and temporal attributes) such as, topology, metric, time range, and alignment; and b) prioritizes such attributes on the basis of a *measurement analysis objective*. In order to strengthen the measurement data “veracity” information, we propose a novel data sanity checking scheme. The measurement data sanity scores are then extended to use a novel Bayesian Inference-driven historical domain/community reputation scheme that acts as “confidence indicators” to help the operators and users in relevant long-term measurement subscriptions. The data sanity together with domain reputation ultimately provides essential meta-information on top of measurement recommendation for the operators/users to take informed decisions.

We use real long-term and short-term perfSONAR traces to show: a) how our scheme recommends most pertinent traces from a large set of available traces from Department of Energy (DOE) Lab sites’ archives based on similarity; and b) how different DOE perfSONAR end-points demonstrate varied characteristics of domain reputation in terms of measurement data quality over a one-year time period. Finally, using an exemplar *spatial measurement analysis* case study with synthetic data that closely mimics real perfSONAR traces, we show: a) how network anomaly event detection using traces recommended by our proposed content filter ensures accurate detection; and b) the relative advantages of content filter based recommendation in terms of false alarms over recommendation approaches that consider partial trace features for filtering, and greedy recommendation approaches based on random trace selection. Overall, the evaluation results demonstrate how our proposed measurement recommendation scheme will enable network operators and application users to intelligently use subscribed community traces for diagnosing network events that impact Big Data applications.

The rest of the paper is organized as follows: Section II discusses the related works in this problem area. Section III motivates the problem in greater depth. Section IV introduces our content-based filtering approach. Section V proposes our Bayesian Inference-based domain reputation scheme. Section VI discusses the evaluation and results using real and synthetic data. Section VII concludes the paper.

II. RELATED WORKS

Content-based social and community infrastructures have been proposed and developed for different areas of computing. However, in relation to measurement frameworks, such works are limited. In [7], the authors propose a tagging and trust mechanism in social networks based on users and their contents that is similar to our research of building a reputation scheme for perfSONAR domains within scientific communities. In [9], the authors propose a Social Cloud Computing environment, outlining various aspects of a social

cloud with a resource/service sharing framework that utilizes relationships established between the members of a social environment. Filter-based (mostly content-based and collaborative) social/community environments have been proposed in different areas of computing. In [10], the authors discuss concepts and applications of collaborative and content-based filtering schemes for social computing environments. In [12], the authors compare the most popular collaborative filtering techniques such as memory-based, model-based, and hybrid, in a recommendation system. Our proposed measurement recommendation scheme although has similarities to existing schemes, it is the first to use content-based filtering for multi-domain network performance monitoring federations.

The perils of using potentially misleading data, and related guidelines to measurement best practices were first highlighted in [3]. Our previous work on using sanitized measurement data for anomaly detection [5] is closest to the work by [16] where an anomaly detection system is developed based on prediction of upper and lower dynamic thresholds of various time-varying data trends. In [17], the authors proposed an overlay fault diagnosis framework with a diagnosis uncertainty reasoning analysis based on evidences similar to our work in [5]. Although our earlier work [5] proposes a measurement data sanitation scheme to quantify the certainty of a detected anomaly event, in this paper we seek to extend such instantaneous data sanity into a historical and quantifiable reputation score of the associated domain(s) and communities. Similar reputation-based trust schemes have long been used by the scientific community for decision making in shared environments. In [8], the authors present a reputation-based trust model for peer-to-peer eCommerce communities. Whereas, in [13], the authors describe a similar scheme to use Bayesian Inference to build reputations for agents in the e-business community. Further, works such as [14] extend such reputation models by introducing an age factor in Bayesian Inference as it is desirable to give greater weight to more recent ratings. Reputation quantification techniques used in all such related works lay the foundation for our proposed novel Bayesian Inference based reputation model for measurement domains that can foster performance tuning of Big Data applications in shared community infrastructures.

III. PROBLEM STATEMENT AND MOTIVATION

Multi-domain data-intensive applications monitor end-to-end network performance by installing perfSONAR MPAs at source and destination domain end points and at strategic locations along the intermediate path. Network operators and application users in charge of such monitoring generally watch for change events in measurement traces. However, often analyzing traces from one such standalone path proves insufficient for effective detection of critical events and subsequent troubleshooting. Thus, they subscribe to measurements from other domains and data-intensive science communities using perfSONAR to collaboratively share measurement data, analyze the shared data, and effectively troubleshoot relevant network events of interest. However, the current perfSONAR framework has limited tools and public services to empower the operators and application users to find the most relevant and valid data that can help effective troubleshooting. Thus, in most cases, operators resort to random trace subscriptions

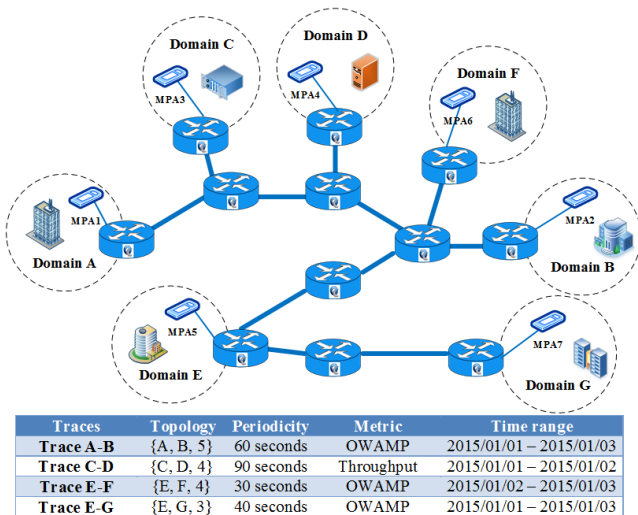


Fig. 1: Network measurement attributes for different traces and the non-triviality of comparison

that could lead to erroneous detection and ineffective as well as time-consuming troubleshooting efforts.

Fig. 1 shows one such scenario where a network operator seeks to perform certain *measurement analysis*, such as, detecting correlated network anomaly events on a network path between domains A and B that involves a data-intensive application flow. In order to achieve this, the operator has collected Trace/s A-B (that we call ‘target’) with specific features such as, network topology, periodicity, measurement metric and time stamp, as shown in Fig. 1. However, in most cases, traces involving one such standalone path, metric and time stamp prove to be insufficient in solving critical network anomaly events that impact an application flow. Thus the operator has to rely on other candidate traces, i.e., Traces C-D, E-F, and E-G that may have completely varied attributes set involving different domains that may or may not have correlation with the target Trace A-B. In such a scenario, it is non-trivial for the operator to decide which among these traces will be more useful to successfully and accurately perform the specific *measurement analysis objective*.

The specific questions we seek to answer in this paper to aid the network operators and application users in the scenario showed in Fig. 1 are: a) Which are the most relevant among the candidates traces (C-D, E-F and E-G) in terms of the specific *measurement analysis objective* that the operator wants to perform? b) How does the operator ascertain the quality of the most relevant traces (among Traces C-D, E-F and E-G) in terms of measurement calibration and data validity so as to use only the pertinent traces for more accurate detection? c) Which candidate domains’ (among C, D, E, F, and G) measurement data should be subscribed in future that will not only be relevant to troubleshoot specific network events of interest but also are most likely to provide reliable quality data to ensure effective troubleshooting? Next, we discuss the “content filter” and “domain reputation” based measurement recommendation scheme that assists the network operators and application users to answer the above questions.

IV. CONTENT-BASED FILTERING MEASUREMENT RECOMMENDATION

In this section, we propose our filter-based measurements recommendation scheme. Our filtering approach is derived from the concepts of content-based filtering techniques used in many recommendation systems, especially by online retail enterprises, such as Amazon, eBay. Content-based filtering, also referred to as cognitive filtering, recommends items based on a comparison between the contents/features of the items and users’ profiles or interests.

Our proposed content filter based recommendation scheme filters and ranks most relevant measurement traces from a pool of traces based on measurement traces’ *attribute similarity* and *measurement analysis objective*. We argue that for any broad type of correlation analysis objective, i.e., spatial or temporal, there are four important factors that define time-series measurement traces’ attributes. They are: 1) topology, i.e., the path taken by the measurement probe packets; 2) metric, i.e., the measurement tool (one-way delay, throughput) used to collect the samples; 3) time range, i.e., the time stamp of the time-series measurements; 4) alignment, i.e., the relative positions of the measurement sampling time stamp instances. For each of these factors, we will quantify the relative similarity between target trace and candidate traces and then create an overall similarity score from individual factor similarities. The overall similarity scores will help the network operators to rank the candidate traces in the order of their relevance.

A. Measurements Topology Similarity

Each perfSONAR measurement trace contains results from an end-to-end probing with synthetic packets that start at the source MPA, and traverse a set of hops to reach a destination MPA. Thus, the topology attribute of each perfSONAR trace can be represented as:

$$topo := (source, destination, hops) \quad (1)$$

where *hops* is the set to represent the intermediate nodes/routers in the topology.

This topology information is very important for correlating measurements traces because - with the similarity between traces’ topologies, probability of common network events of interest increases. In our previous experiences with correlated anomaly detection in perfSONAR traces [2], we observed that traces with common hops are more useful than others to establish correlation between anomaly events and to diagnose the root causes more accurately. Therefore, we express the topology similarity $topo_simi_{i,j}$ between target trace *i* and candidate trace *j* as:

$$topo_simi_{i,j} = \frac{topo_i \cap topo_j}{topo_i \cup topo_j} \quad (2)$$

B. Measurements Metric Similarity

The measurement *metric* indicates the network performance measurement tool used for monitoring, such as Ping, OWAMP, BWCTL, etc. The measurement metrics similarity between traces is of importance based on the type of measurement analysis sought. For example, in spatial analysis, such as correlated anomaly detection, operators tend to use only homogenous traces in terms of metrics to diagnose a root cause location.

Whereas, in temporal analysis, operators rely on different metrics such as, OWAMP, traceroute to detect events in a particular time window. Thus, we express measurement metric similarity $metric_simi_{i,j}$ between traces i and j as a boolean representation:

$$metric_simi_{i,j} = \begin{cases} 1, (m_i = m_j) \\ 0, (m_i \neq m_j) \end{cases} \quad (3)$$

C. Measurements Time Range Similarity

The measurement *time range* of a trace is one of the most important measurement attributes for temporal analysis when the duration of the traces becomes critical to detect a time-specific network event. Thus, if two traces' durations are not aligned temporally, their time range similarity should be equal to zero. Therefore, we express the time range similarity $range_simi_{i,j}$ between traces i and j as:

$$range_simi_{i,j} = \frac{r_i \cap r_j}{r_i} \quad (4)$$

D. Measurements Alignment Similarity

Measurement alignment, i.e., the relative positions of measurement sample instances is also significant for correlation analysis with multiple dimension time-series measurements. Samples that are closely aligned are easier to correlate and have better chances of accurate detection of network events upon analysis. Although, such alignment is closely related to the measurement sampling periodicity, the periodicity itself cannot be used as a factor in similarity matching. This is because, in practice it is difficult to find perfectly periodic samples; sample periods always fluctuate around a mean value. Thus, relative alignment of sample instances between two traces is a better metric to quantify their relative similarity. In an illustrative example shown in Fig. 2, we show one target and three candidate traces with different periods and sampling patterns. As far as the similarity is concerned, candidate trace 1 is best aligned to the target trace as the mean relative displacement between the trace 1 and the target trace is minimum.




Fig. 2: Measurements alignment illustration between traces with varied periodicity and non-aligned sample time stamp instances

Thus, for generic quantification of alignment between measurement traces, we define an alignment displacement metric d that denotes the mean relative distance between sample instances of target trace i and candidate trace j . The metric d is expressed as:

$$d_{i,j} = \sum_T |ts_i - \hat{ts}_j| \quad (5)$$

where ts_i denotes target trace time stamp, \hat{ts}_j denotes the candidate time stamp closest to the target trace time stamp ts_i , and T denotes the number of such time stamps in the

target trace. The range of metric d varies between $[0, +\infty)$ with smaller value indicating better alignment between traces. In order to normalize d and be consistent with other similarity factors, we transform the range of metric d between $[0, 1]$ using min-max normalization method, which can be expressed as:

$$align_simi_{i,j} = \frac{max_{d_{i,j}} - d_{i,j}}{max_{d_{i,j}} - min_{d_{i,j}}} \quad (6)$$

where $max_{d_{i,j}}$ and $min_{d_{i,j}}$ are the maximum and minimum values of $d_{i,j}$.

E. Overall Similarity Scores and Analysis Objective

The overall measurements similarity is expressed as a weighted product of the aforementioned four similarity factors:

$$overall_simi_{i,j} = \sum_k w_k * factor_simi_{i,j}^k \quad (7)$$

where $factor_simi_{i,j}^k$ denotes each of the aforementioned similarity factors and w_k denotes their respective weights for the overall measurements similarity score. The values of the weights depend on the relative importance of these attributes to achieve different *measurement analysis objectives*. For example, topology information is very important in spatial analysis, such as correlated anomaly detection, thus making weight w_k for $topo_simi_{i,j}$ greater than other weights. In order to simplify the weights but at the same time to have a more comprehensive list of analysis objective scenarios, we design a decision tree of different generic measurement analysis objectives and corresponding relative importance of attributes' weights (w_t , w_m , w_r , and w_a respectively for each of the similarity factors), as shown in Fig. 3.

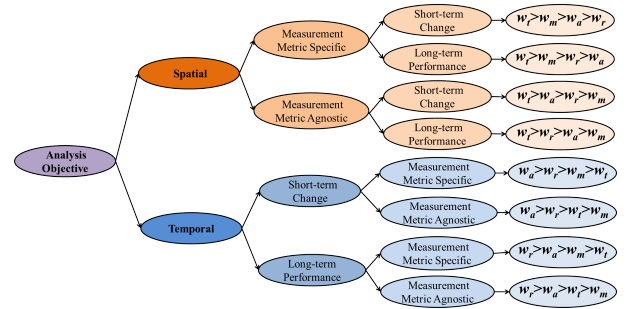


Fig. 3: Decision tree for different measurement analysis objectives and the corresponding relative measurement attributes' weights

Broadly, we divide the entire analysis objective space into temporal and spatial analysis. In temporal analysis, the weights for temporal factors, such as *time range* and *alignment* are more important than spatial factors such as *topology*. These categories are then further divided on the basis of relevance of measurement metrics, and short and long term analysis. The leaf nodes of the tree denote the relative importance of attributes' weights. We argue that most types of the measurement analysis required for end-to-end data-intensive application performance monitoring fall under one of these subcategories.

For example, one of the most common measurement analysis for end-to-end data-intensive application management is the correlated anomaly event detection [2], which involves a spatial analysis and falls under 'Analysis Objective' → 'Spatial' → 'Measurement Metric Specific' → 'Short-term Change'. So, the relative weights should be $w_t > w_m > w_a > w_r$.

V. DOMAIN REPUTATION

The concept of reputation is closely linked to trustworthiness; an entity’s reputation is generally a subjective proof of its historical actions and in most cases a measure of expectations of future behavior. Reputation schemes are widely used in e-commerce systems to build long-term beliefs for agents and encourage good behavior. The main objective of proposing a domain reputation scheme is to generate domain-centric expectations for operators when they subscribe to perfSONAR measurement data from different domains. Such a domain reputation scheme can also encourage good measurement practices (e.g., sharing of calibrated measurement tool data) among domains. Such a domain reputation metric and ensuing expectation will add a new dimension to the measurement recommendation system to help operators and application users make more educated collaborative decisions in multi-domain infrastructure performance management.

Reputation of any domain is a function of the quality of measurement data generated from that domain. In our previous work [5], in order to ascertain what features of a sample set of measurement data qualify them as ‘good’, we collected a considerable amount of perfSONAR one-way delay (OWAMP) traces for different paths and different time periods. In any random collection that are publicly accessible, we observed that some measurements exhibit non-periodic sampling pattern, i.e., either too dense or too sparse, and some are invalid due to faulty clock synchronization between measurement servers or data corruption (e.g., negative one-way delay values). Therefore, we argue that it is of paramount importance that subject to a given *measurement analysis objective*, the collected perfSONAR samples should exhibit some desired characteristics, which for spatial analysis with OWAMP are: *sampling pattern* and *data validity*. Thus, in our earlier work [5], we have defined the sanity score of any trace with path (source destination pair) i as:

$$s_i = \frac{N_i - (N_i - n_i^{majority}) - (N_i - n_i^{valid})}{N_i} \quad (8)$$

where N_i denotes the number of measurement samples in path i , n_i^{valid} denotes the number of valid data (without faulty clock synchronization and/or data corruption) samples in path i , and $n_i^{majority}$ denotes the number of samples showing consist sampling pattern.

In this paper, we employ Bayesian Inference to translate such data sanity scores into a domain reputation. Through Bayesian Inference, new or an updated reputation score (i.e., posteriori) of an entity can computed by combining the old/previous reputation score (i.e., priori) with a new belief. In order to translate sanity scores into domain reputation, we first discretize the measurement data sanity scores into data sanity ratings of a particular domain using boolean variables such as ‘Good’ (variable x) and ‘Bad’ (variable y), and some sanity threshold ϵ . The value of ϵ is a measure of the degree of conservativeness of the reputation scheme that is kept constant for the entire system. The value of ϵ can be set based on the distribution of measurements’ sanity scores in the system. If the average sanity score of measurement data in the system is very high, ϵ value is kept high to differentiate between good and bad measurements, and vice versa. Usually for all practical purposes, ϵ value is around $[\mu + \sigma, \mu + 2 * \sigma]$.

$$x = |i| \quad \forall s_i \geq \epsilon; \quad y = |i| \quad \forall s_i < \epsilon \quad (9)$$

Therefore, at any given time t , the measurement data sanity rating of any domain is represented as $\rho^t = [x, y]^t$. Now if there are T such discrete data sanity ratings collected over a period of time, then the overall data sanity rating after T collection is given as $\rho^T = [x, y]^T$, where x^T and y^T are expressed as:

$$x^T = \sum_{t=1}^T \lambda^{T-t} x^t \quad \text{and} \quad y^T = \sum_{t=1}^T \lambda^{T-t} y^t \quad (10)$$

where $0 \leq \lambda \leq 1$ is called the ‘forgetting factor’ and keeps the recent history of data sanity rating more relevant in reputation calculation than ancient history. The value of λ represents how forgetful a system is, with $\lambda = 1$ means the system forgets nothing. Thus, this value is user opinion dependent, such that, if the user thinks the historical reputation is also very important, this value should be close to 1; however, if the user thinks current reputations are more important, the value should be close to 0.

Now after collecting T such discrete data sanity ratings, the reputation of the domain responsible is expressed as a posterior expectation of beta distribution of ρ^T and is represented as:

$$R^T = E[\text{beta}(\rho^T)] \\ = \frac{x^T + 1}{x^T + y^T + 2} \quad (11)$$

VI. EVALUATION

We evaluate the performance of our proposed measurement recommendation scheme using real and synthetic perfSONAR measurements. Real perfSONAR measurement traces are used to demonstrate the working of the recommendation scheme and domain reputation establishment. Whereas, the synthetic data is used to showcase the effectiveness of the recommendation scheme in terms of using the recommended samples for accurate detection of correlated network anomaly events.

A. Recommendation scheme functionality with real traces

We collect hundreds of perfSONAR traces from DOE labs and ESnet sites with perfSONAR end-points with different measurement attributes as inputs to our proposed measurement recommendation scheme. In Table I, we show the attributes for only a small subset of collected samples. In this subset, we have kept the trace BNL↔FNAL as the target trace (in blue color font) and the rest as candidate traces. Applying our proposed recommendation scheme, we seek to find the most relevant traces for two separate measurement analysis objectives: a) Temporal analysis for correlated anomaly event detection [2] where complete topology information is not available (which is often the case for public perfSONAR datasets); and b) Spatial analysis for topology-aware correlated anomaly event detection [15].

TABLE I: Real perfSONAR traces’ attributes description

Trace Name	Metric	Periodicity	Time Range
bnl ↔ fnal	One-way delay	[58, 62]	2015-09-01 00:00:57 ↔ 2015-09-09 23:58:59
lbl ↔ orn1	One-way delay	[51, 68]	2015-09-03 05:39:18 ↔ 2015-09-03 22:19:49
aofa ↔ bost	One-way delay	[51, 161]	2015-09-01 00:00:35 ↔ 2015-09-01 05:35:22
bnl ↔ bois	One-way delay	[57, 63]	2015-09-01 00:00:20 ↔ 2015-09-09 23:59:46
sacr ↔ bois	One-way delay	[53, 150]	2015-09-01 00:00:21 ↔ 2015-09-01 16:42:37
bnl ↔ bost	One-way delay	[53, 67]	2015-09-01 00:00:53 ↔ 2015-09-01 05:35:20
hous ↔ srs	One-way delay	[53, 69]	2015-09-05 05:31:18 ↔ 2015-09-05 22:08:47
bnl ↔ lbl	One-way delay	[61, 66]	2015-09-01 00:00:17 ↔ 2015-09-09 23:59:13
newy ↔ sacr	Throughput	[54, 67]	2015-09-01 00:00:52 ↔ 2015-09-09 23:59:15
anl ↔ newy	Throughput	[52, 64]	2015-09-01 00:01:00 ↔ 2015-09-09 23:59:47
bnl ↔ nash	Throughput	[57, 71]	2015-09-03 05:58:28 ↔ 2015-09-03 22:38:15
denv ↔ fnal	Throughput	[58, 67]	2015-09-01 00:00:51 ↔ 2015-09-09 23:59:58

Table II shows the measurement attributes similarity scores for the candidate measurement traces with the target trace described in Table I. The score evaluations follow the scheme described in Section IV and due to the varied attributes of the collected traces, we observe that the similarity scores of candidate traces based on different attributes can vary by a considerable margin.

TABLE II: Measurement attributes similarity score description

Trace Name	Topology Similarity	Metric Similarity	Alignment Similarity	Time Range Similarity
lbl ↔ orn1	0	1	0.456	0.077
aofa ↔ bost	0.077	1	0	0.026
bnl ↔ bois	0.385	1	0.999	1
sacr ↔ bois	0	1	0.103	0.077
bnl ↔ bost	0.4	1	0	0.026
hous ↔ srs	0	1	0.767	0.53
bnl ↔ lbl	0.333	1	0.999	1
newy ↔ sacr	0.125	0	0.999	1
anl ↔ newy	0	0	1	0.999
bnl ↔ nash	0.143	0	0.979	0.862
denv ↔ fnal	0.154	0	0.999	1




Fig. 4: Historical reputation characteristics comparison of 3 exemplar DOE lab sites with real perfSONAR traces over one year period

In Fig. 4, we show the historical reputation characteristic comparison of 3 exemplar DOE lab sites based on one year (Oct 2014 - Oct 2015) traces’ data sanity scores using the scheme discussed in Section V. We observe that although these are well known and seemingly reputed perfSONAR end-points within the the data-intensive application communities deploying perfSONAR, not all the sites produce trustworthy data at all times. Thus, we establish the need for a domain’s reputation as a key factor in subscribing to that domain’s measurement data for accurate detection and effective troubleshooting. For example, according to Fig. 4, subscribing to data from domain ‘NEWY and ‘ATLA’ is likely to yield data with high veracity; whereas, subscribing to STAR data may not always lead to accurate correlated anomaly event detection and diagnosis.

Finally, we apply the relative weights of the measurement attributes (as shown in Fig. 3) on the similarity scores based on the two monitoring objectives: temporal and spatial. For temporal analysis, we focus on detecting correlated anomaly events in time series measurements where measurement metric needs to be of similar type. Hence, we follow the path: ‘Analysis Objective’ → ‘Temporal’ → ‘Short-term Change’ → ‘Measurements Metric Specific’, and assign the weights according to the rule $w_a > w_r > w_m > w_t$. Whereas for spatial analysis, focus is on using measurements topological information to drill down the location of events. Hence, it follows the path: ‘Analysis Objective’ → ‘Spatial’ → ‘Measurements Metric Specific’ → ‘Short-term Change’, with final relative weights following the rule $w_t > w_m > w_a > w_r$.

The final recommendation outcomes and corresponding ranking of a subset of candidate traces are shown in Table III

along with the traces’ data sanity scores, and corresponding source and destination domains’ reputation scores. Table III is a snapshot of the actual manifestation of our proposed recommendation scheme to assist the operators and application users to better gauge the relevance and veracity of collected samples. For example, the Trace BNL↔BOIS (in blue font color as shown in Table III) is the best choice among the candidates (1st ranked) in terms of similarity and high instantaneous data sanity score (0.993). Whereas, Trace BNL↔LBL (in red font color), although being 2nd ranked for both temporal and spatial analysis, may not be a good choice for candidacy as the low sanity score (0.642) suggests sub-par data quality which can be attributed to low reputation (0.611) of the destination domain (LBL). Thus, operators should be advised to use Trace DENV↔FNAL (in green font color) over Trace BNL↔LBL for temporal analysis as the data quality of the former is much better (0.972) in spite of having a slightly lower similarity (0.765). However, DENV↔FNAL will not be a significantly better choice over BNL↔LBL for spatial analysis due to the former’s very low similarity score (0.327).

B. Recommendation scheme effectiveness with synthetic traces

In order to examine the effectiveness of our proposed recommendation scheme in accurately identifying anomaly events upon analysis, we perform experiments with synthetic perfSONAR data. We randomly generate a one-week dataset, and inject different number of correlated and uncorrelated anomaly events to examine the detection accuracy. The synthetic data is carefully generated to closely mimic the actual perfSONAR OWAMP measurement traces. In order to inject correlated anomaly events, we generate 100 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 varies from 0.1%-1% of the trace sample population. The magnitudes of anomaly events vary from 10% - 60% over normal measurements with higher magnitudes causing sharper spikes.

In the first experiment, we use the traces recommended by our scheme to detect network anomaly events using an exemplary temporal correlation analysis. For this experiment, we use different number of recommendations and see whether analysis with such recommendations can successfully identify the correlated/uncorrelated anomaly events that we injected. In Figs. 5(a), 5(b), and 5(c), we show the accuracy of such anomaly event detection with our scheme recommending 2, 10 and all 100 traces, respectively. We observe that in this particular scenario, 10 recommendations accurately detect all the anomaly events with no false alarms. Whereas analysis with only 2 recommendations lack the necessary data to establish correlation, thus causing false alarms. Further, all 100 recommendations suffer from too much noise in anomaly detection caused by undesired traces resulting in false alarms. In Fig. 5(d), we showed the nature of detection accuracy with the number of recommended traces exhibiting an inflection within 10 - 40 recommendations range. Thus, we argue that there exists: i) an optimal number of recommendations (in this case 10) for accurate detection, and ii) an inflection point (in this case 40) beyond which too many traces contribute to high levels of noise resulting in false alarms.

TABLE III: Data sanity score and domain reputation results for selected traces used in exemplar analysis case study

Trace Name	Data Sanity Score	Temporal Correlation Analysis		Spatial Correlation Analysis		Domain Reputation	
		Overall Similarity Score	Ranking	Overall Similarity Score	Ranking	Source Reputation	Destination Reputation
bnl ↔ bois	0.993	0.938	1	0.938	1	0.983	0.984
bnl ↔ lbi	0.642	0.933	2	0.933	2	0.985	0.611
denv ↔ final	0.972	0.765	3	0.327	5	0.991	0.984
newy ↔ sacr	0.982	0.762	4	0.312	7	0.979	0.986
anl ↔ newy	0.984	0.221	5	0.273	10	0.987	0.986
bnl ↔ boost	0.953	0.197	10	0.453	3	0.983	0.964
hous ↔ srs	0.976	0.666	7	0.418	4	0.934	0.986

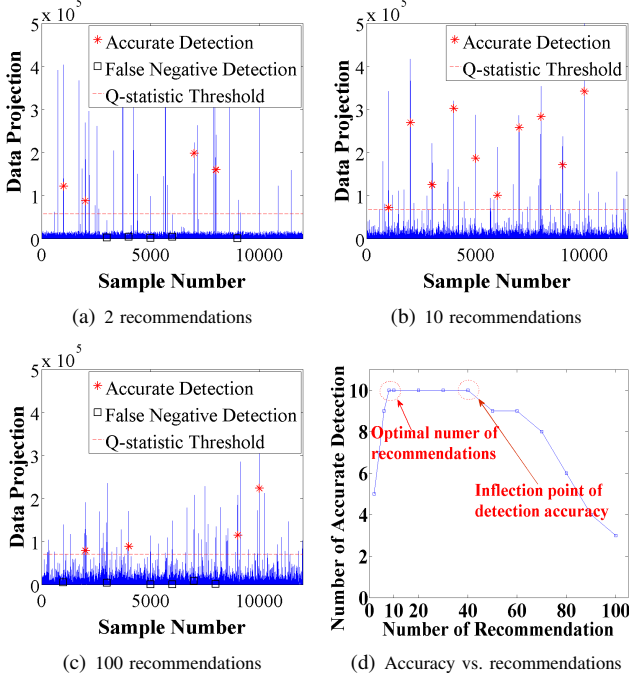


Fig. 5: Accuracy of correlated anomaly detection in terms of false alarms with varying number of recommended traces




Fig. 6: False Alarm Rate comparison among the three different schemes evaluated

Fig. 6 shows the benefits of our content filter based recommendation scheme over the greedy recommendation approach (random selection), and recommendation strategy with filtering based on partial measurement features, such as temporal aspects (time range) of the traces. For this experiment, we vary the density of anomaly events and recommend equal number of traces for each approach being compared. We see that on an average the content filter performs consistently better than the greedy and temporal filter based approaches. It is interesting to observe that for a very small density of anomaly events, the false alarm rate is higher for all approaches. This is because, too few anomaly events with minimal anomalous features are difficult to detect and are neither related to the number of recommendations nor the filtering approaches.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we established the need of a measurement recommendation for network operators and users to effectively troubleshoot bottlenecks affecting Big Data applications. Using a content-based filter, we proposed a measurement recommendation scheme that filters and ranks best relevant traces from a pool of candidate traces. The recommendation scheme was complimented by a Bayesian Inference based domain reputation calculation scheme that indicates the trustworthiness of the collected samples among the involved domains. Using real and synthetic perfSONAR short-term and long-term traces, we showed how our content filter enables operators to intelligently use less but relevant measurement samples to accurately detect and diagnose performance bottleneck causing network events.

In future, we will extend our measurement recommendation scheme to incorporate the principles of collaborative filtering. Such a collaborative filtering approach will allow the operators and application users to collectively troubleshoot common network problems through similarity in multiple user perspectives on top of similarity in measurement traces. Our work will thus help data-intensive application communities to better socialize around measurements and achieve expected performance.

REFERENCES

- [1] A. Hanemann, J. Boote, E. Boyd, J. Durand, L. Kudrimoti, 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] Y. Zhang, P. Calyam, S. Debroy, M. Sridharan, "PCA-based network-wide correlated anomaly event detection and diagnosis," *Proc. of DRCN*, 2015.
- [3] V. Paxson, "Strategies for sound internet measurement", *Proc. of ACM SIGCOMM IMC*, 2004.
- [4] P. Calyam, J. Pu, W. Mandrawa, A. Krishnamurthy, "OnTimeDetect: Dynamic Network Anomaly Notification in perfSONAR Deployments", *Proc. of MAS-COTS*, 2010.
- [5] Y. Zhang, S. Debroy, P. Calyam, "Network-wide Anomaly Event Detection and Diagnosis with perfSONAR", *IEEE TNSM*, 2016.
- [6] P. Calyam, L. Kumarasamy, C. Lee, F. Ozguner, "Ontology-based Semantic Priority Scheduling for Multi-domain Active Measurements", *Springer JNSM*, 2014.
- [7] I. Ivanov, P. Vajda, J. Lee, T. Ebrahimi, "In Tags We Trust: Trust modeling in social tagging of multimedia content", *IEEE Signal Processing Magazine*, 2012.
- [8] L. Xiong, L. Ling, "A Reputation-Based Trust Model for Peer-to-Peer eCommerce Communities", *Proc. of the IEEE Conference on E-Commerce*, 2003.
- [9] K. Chard, K. Bubendorfer, S. Caton, O. Rana, "Social cloud computing: A vision for socially motivated resource sharing", *IEEE TSC*, 2012.
- [10] M. Tavakolifard, K. Almeroth, "Social Computing: An Intersection of Recommender Systems, Trust/Reputation Systems, and Social Networks", *IEEE Network Magazine*, July 2012.
- [11] E. Blanton, S. Fahmy, G.N. Frederickson, "On the Utility of Inference Mechanisms," *Proc. of ICDCS*, 2009.
- [12] X. Su, T. Khoshgoftaar, "A survey of collaborative filtering techniques". *Adv. in Artif. Intell. Article 4*, 2009.
- [13] L. Mui, M. Mohtashemi, A. Halberstadt, "A computational model of trust and reputation," *Proc. of HICSS*, 2002.
- [14] A. Jsang, R. Ismail, "The Beta Reputation System", *Proc. of the 15th Bled Electronic Commerce Conf.*, 2002.
- [15] P. Calyam, M. Dhanapalan, M. Sridharan, A. Krishnamurthy, R. Ramnath, "Topology-Aware Correlated Network Anomaly Event Detection and Diagnosis", *Springer JNSM*, 2013.
- [16] M. Marvasti, A. Poghosyan, A. Harutyunyan, N. Grigoryan, "An Enterprise Dynamic Thresholding System", *Proc. of ICAC*, 2014.
- [17] Y. Tang, E. Al-Shaer, K. Joshi, "Reasoning under Uncertainty for Overlay Fault Diagnosis", *IEEE TNSM*, 2012.