On the Use of a Measurement Correlation Service for Measurement Federations

Jéferson Campos Nobre*, Leandro Lisboa Penz†, Muriel Figueredo Franco ‡, Lisandro Zambenedetti Granville‡
* University of Vale do Rio dos Sinos, Brazil
† Arista Networks, Ireland
‡ Federal University of Rio Grande do Sul, Brazil
jcnobre@unisinos.br, lpenz@lpenz.org, {mffranco, granville}@inf.ufrgs.br

***Abstract*—The diversity of services that operate in the Internet has increased significantly in the last years. Performance problems in these services cause important financial losses. To ensure that these problems do not occur, service levels need to be monitored. One of the main techniques for such monitoring involves the use of active measurement mechanisms. However, these mechanisms are expensive in terms of resources consumption due to the activation of measurement sessions. Measurement sessions usually cover only a fraction of what could be measured, which leads to service level problems being missed. Measurement federations can help network administrators in different tasks, such as controlling the activation of active measurement sessions. In this context, measurement correlation can be deployed in order to improve this control in such federations. The main contribution of the present work is the proposal of a data transformation service that provides measurement correlation. This service is used to enable cooperation features in measurement federations, while decreasing resource consumption. Besides that, statistical tests that can be used to compose such correlation are presented. The proposed solution is evaluated using an active measurement dataset from the Brazilian National Research and Education Network (*Rede Nacional de Ensino e Pesquisa* - RNP). Our results provide insights regarding measurement correlation from federated measurement points and can be used for the design of better application to control active measurement sessions.**

I. Introduction

Significant progress has been made in increasing the capacity and accessibility of networking infrastructures over the last decades. Alongside, the diversity of applications and services that operate in such infrastructures has increased as well. Unfortunately, such progress is not always directly translated into improved performance of all applications and users. Performance issues can vary because of traffic congestion to packet loss caused by intermediate link failures. Service level monitoring tools must be scalable and capable of detecting performance issues in a timely and efficient manner. Considering such tools, measurement mechanisms are one of the most used.

Service level monitoring can be performed using either passive or active mechanisms. Network conditions are monitored in a non-intrusive way with passive mechanisms (*e.g.*, IP Flow Information EXport - IPFIX [1]), since traffic is not created by the monitoring process itself. Active measurements (*e.g.*, One-Way Active Measurement Protocol - OWAMP [2]), on the other hand, are intrusive because they inject monitoring traffic on the network infrastructure to deliver performance metrics. The service level monitoring is usually performed through active mechanisms because it usually presents better accuracy and privacy features than passive ones. On the basis of active mechanisms, a larger monitoring coverage increases the consumed resources (*e.g.*, CPU cycles and memory footprint) and network load due to the amount of sessions activated to measure a destination. The best practice to tackle this problem is to define the measurement sessions considering the expertise and knowledge of the network administrator. For this, measurement federations can be a valuable tool.

Measurement federations could help network operators troubleshoot perceived abnormalities as well as improve network middleware regarding service level issues. Examples of those federations are the ones implemented using the PERFormance Service Oriented Network monitoring ARchitecture (perfSONAR) toolkit [3]. In this context, active measurement mechanisms are one of the most important measurement tools in such federations. However, the sole employment of a measurement federation does not improve the capabilities of these mechanisms in important aspects, such as scalability and efficiency. The current best practice regarding active measurements mechanisms, for example, usually covers only a fraction of the network flows that should be observed in order to save resources. This can lead to performance problems being missed. Besides, this practice is labor-intensive for the network administrator and inefficient considering highly dynamic network infrastructures, since administrator reactions are necessary to reconfigure the active measurement mechanisms for monitoring different network destinations. On the basis of these issues, novel approaches are needed to improve the operation of active measurement mechanisms in measurement federations.

Despite the fact that measurement federations are composed of programmable nodes, control loops are not used for measurement sessions activation. It makes difficult the employment of active mechanisms. However, approaches based on measurement correlation have been effective to enable control loops for measurement session activation [4]. Based on this, in this paper, we present a solution to enable the measurement correlation on federations. This solution extends some of our preliminary concepts and results [5]. The main contributions of this work are: i) the support for improved measurement coverage considering federations with significant measurement correlation among participating nodes; ii) the definition of a measurement correlation service considering a well-known measurement federation; and iii) the analysis of statistical tests to perform measurement correlation. The proposed solution was evaluated using a dataset from *MonIPÊ*, a perfSONAR-based measurement federation related to the IPÊ Network. This network is supported by the Brazilian National Research and Education Network (*Rede Nacional de Ensino e Pesquisa* - RNP).

The remaining of this paper is organized as follows. In Section II, we describe the background and related work on active measurement mechanisms and measurement federations. In Section III, we detail our proposal to introduce destination correlation in measurement federations. In Section IV, we present an experimental evaluation. Finally, in Section V, we discuss conclusions and future work.

II. Background and Related Work

In this section, we first cover one of the most prominent active measurement mechanisms, the One-way Active Measurement Protocol (OWAMP), which is designed to measure one-way latency, packet loss, packet duplication, and jitter on an end-to-end basis. After that, some widely known measurement federations are presented. In addition, we detail the perfSONAR, a network measurement toolkit that provides federated measurements. Finally, we present the related work.

A. One-way Active Measurement Protocol

The IETF IP Performance Metrics (IPPM) Working Group (WG) proposed open mechanisms that permit the exchange of packets to collect metrics for one-way (One-way Active Measurement Protocol - OWAMP) [2] packet delay and loss across Internet paths in an interoperable manner. OWAMP actually consists of two inter-related protocols: OWAMP-Control and OWAMP-Test. OWAMP-Control is layered over TCP and used to initiate, start, and stop test sessions and fetch their results. OWAMP-Test, on the other hand, uses UDP to exchange test packets between two measurement nodes. Together, these protocols enable the standardized execution of one-way measurements.

The OWAMP tool is an implementation of OWAMP, developed by the Internet2 Consortium. This tool comprises several applications for the execution of active measurement tasks, such as *owping* and *owstats*. One of these applications is *powstream*, which creates a continuous stream of one-way latency tests by aligning together multiple OWAMP test sessions. This application works by contacting a daemon on the remote host. *Powstream* produces data in two formats: raw owamp data files and summary statistics. The data files are the same binary format saved from owping and can be parsed using *owstats*.

B. Measurement Federations

Sophisticated tools are necessary to monitor large and complex networks, since this requires a great deal of effort and expertise from network administrators. Network measurement toolkits can encompass monitoring and diagnostic tools in order to aid administrators on their daily tasks. In this context, measurement federations are defined as a group of integrated management stations running such toolkits.

Measurement federations can aid network administrator teams and increase productivity considering available computational and human resources. The focus of these federations is to provide a measurement infrastructure capable of monitoring network performance and troubleshooting related issues. Some examples of such federations are the *PERFormance Service Oriented Network monitoring ARchitecture (perfSONAR)* and the *Measurement Lab (M-Lab)*.

The perfSONAR, which is key to this paper, is a network measurement toolkit designed to provide federated measurements. It is the result of an international collaboration for network monitoring, composed of several organizations, such as Internet2, Energy Sciences network (ESnet), Indiana University, and the GÉANT Project. There are several perfSONAR instances deployed world wide and many of which are available for open testing. The perfSONAR toolkit includes *OWAMP*, *Bandwidth Test Controller (BWCTL)*, and *nuttcp* tools. Besides that, the toolkit provides an interface for measurement scheduling, storage of data in uniform formats, and scalable methods to retrieve data. Figure 1 depicts the architecture of perfSONAR measurement framework.

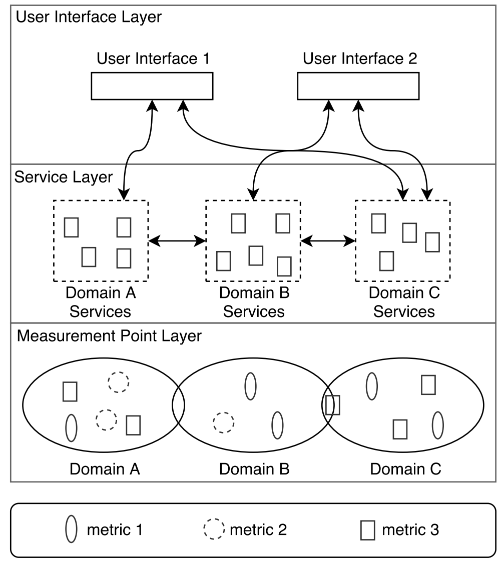

Fig. 1. perfSONAR Measurement Architecture.

There is an inherent human and computational cost related to the deployment of active measurement mechanisms, such as OWAMP, and their continuous operation, *i.e.*, the management of measurement sessions. This cost is strongly influenced by the size and complexity of the network infrastructure. Even using measurement federations, these costs decrease the measurement coverage, because it is usually not possible to monitor all network flows. In this context, solutions to improve the control of measurement mechanisms as well as increase their efficiency are vital.

C. Related Work

The use of measurement correlation to tackle the problem of SLA monitoring was investigated in the context of some research initiatives over the past years. Some of these initiatives are discussed as follows.

Abdelkefi et al. [6] proposed an approach called "Service-quality Characterization of Internet-path" (SCI) to detect abrupt changes in network metrics using active measurements. SCI uses delay and loss measurements collected from vantage points at the two ends of an Internet path. These measurements are transformed into performance signals, in which abrupt changes are detected using Principal Component Pursuit (PCP). The detected abrupt changes in aggregate delay and average delay, together with the loss information, are further mapped as different types of events causing degradations and failures in the service within the path. The authors evaluated on the dataset collected from a real operational environment.

Zhang et al. [7] investigated 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. Such recommendations are complimented with Bayesian Inference-based domain reputation meta-information to strengthen the veracity information of the recommended traces. 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. However, the authors performed the evaluation only using synthetic traces.

III. Measurement Correlation in Federations

Novel techniques in the area of monitoring and troubleshooting have been proposed to tackle resource consumption issues. One of the proposed concepts is the use of correlated peers to enable embedded and collaborative active SLA monitoring [4]. Network devices are correlated peers when they have similar measurement results considering specific metrics. The similarity is quantified through destination scores [4].

Destination scores [4] can be used to provide correlation features in measurement federations ('ds' in Equation 1). For example, specific parts of a network may require more attention due to unexpected events. Thus, the moving average distance of past measurement results to Service Level Objective (SLO) for the destinations can increase the probability of measuring such destinations. As depicted in Equation 1, this average consists of calculating the simple arithmetic mean over values contained in fixed-size subset of a time series (window size, 'ws' in Equation 1) which is permanently shifted as new data arrives (sliding window). Besides that, the time elapsed from the last measurement for a given destination also composes the destination score. This aims at maintaining frequent measurements on destinations.

$$ds = (\frac{1}{ws}\sum_{j=i}^{i+ws-1} a_j) + ts \quad (1)$$

In this section, we present an architecture to deploy measurement correlation considering the perfSONAR toolkit. In addition, we describe some statistical tests that can be used for such correlation.

A. Proposed Architecture

We instantiate our proposed architecture considering the perfSONAR toolkit. This choice is based on the deployed infrastructure on the Brazilian NREN (MonIpê). The key concept of perfSONAR is that each service performs a network performance measurement and analysis function. Some examples of such services are Measurement Point (initiation of performance tests), Measurement Archive (performance monitoring results storing), Transformation Service (operations on data sets), Resource Protector (arbitration on the resources consumption), and Lookup (locating other instances).

In order to support measurement correlation, we describe a Transformation Service (TrS) for Measurement Correlation (MC). The flow and basic operations of our architecture are presented in Fig. 3. TrSs perform operations on data sets that have potential to be stored and replayed. Thus, this service provides a conduit for popular operations. Some roles of TrSs are aggregation, concatenation, correlation, and translation. The TrS MC uses the Measurement Archive (MA) service to collect raw and aggregated OWAMP *powstream* data from the perfSONAR Measurement Points (MPs). Then, statistical functions are employed to compare collected measurements results. The result of a TrS MC operation is the evaluation of

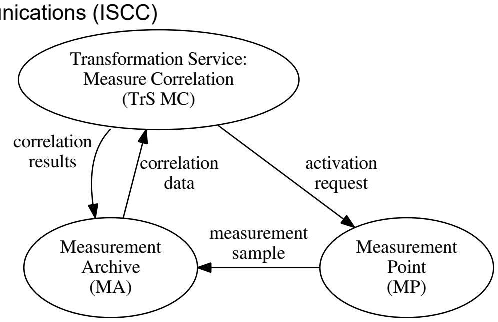

Fig. 2. Architecture flow and basic operations

destination scores. It is not necessary to run the MC TrS on the same machine providing the MA service.

The TrS MC is designed to run statistics over changing datasets. Thus, it can collect measurement results between arbitrary points and considers distributed services, i.e., it can be deployed in different federation nodes regarding MA and MP services. In this context, the TrS MC is compatible with architectures for federated measurement that provide a backend for measurement results storage and a defined interface for data collection and sharing. Besides, the measurement correlation can be further automated if the federation supports registration management and location of distributed services.

B. Statistical Tests for Measurement Correlation

We propose the use of parametric and non-parametric statistics for measurement correlation. Statistical hypotheses about parameters concern the behavior of observable random variables. In parametric tests, the hypothesis is that the data follows a probability distribution model that can be described by a fixed set of parameters. Thus, it is possible to use the parameters to infer the probability of seen and unseen values. Some parametric statistics, for instance, assume a Normal underlying data distribution that has, as parameters, a fixed (maybe unknown) mean and variance. On the other hand, non-parametric statistics assume no probability distribution for the data, and do not use probability distribution parameters. In the present work, we employ parametric and non-parametric statistics which can be divided in two classes: correlation coefficients and distribution tests.

The measurement correlation can be performed using correlation coefficients produced by different tests. Such coefficients are a measure of statistical dependency between 2 random variables, varying between -1 (total negative correlation) and +1 (total positive correlation). The Pearson product-moment correlation coefficient is a parametric correlation measure that represents the linear dependency between two variables. In addition, non-parametric correlation tests, such as the Spearman rank correlation coefficient and the Kendall rank correlation coefficient can be also used.

The distribution tests are used to verify whether samples originate from the same distribution, usually providing a measure that varies between 0 and +1 (statistical significance). The Student's t-test (null hypothesis assumes that the means and variances of two samples are equal) and Welch's t-test (null hypothesis assumes that only the means of two samples are equal). These tests are generalized as Analysis of Variance (ANOVA) tests. Since t-tests consider statistical parameters, they are classified as parametric tests. For exactness, it is required normality of the samples. Besides, t-test requires that the sample variance follows a scaled χ^2 distribution. In order

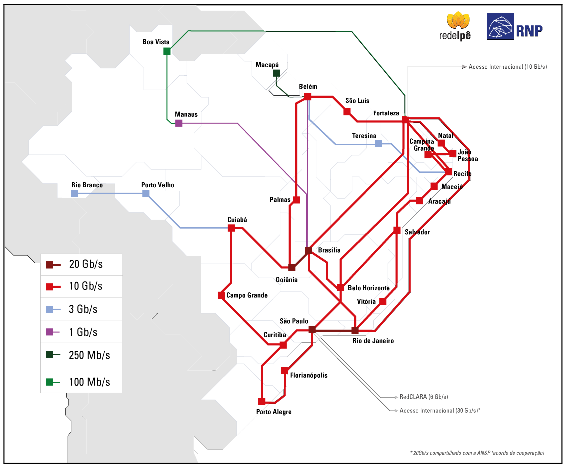

Fig. 3. Logical Topology of the Points of Presence of the IPÊ Network.

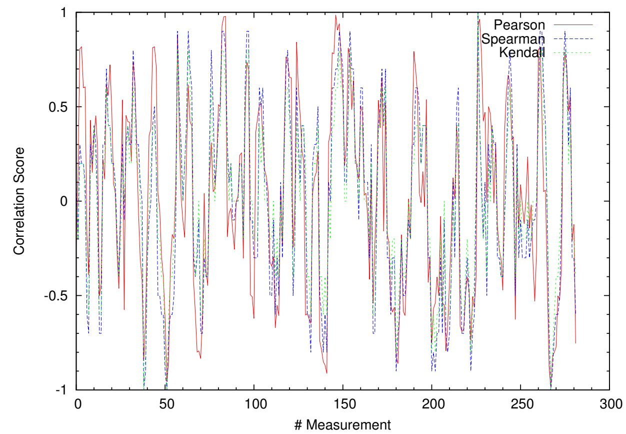

Fig. 4. Correlation Coefficients - Porto Alegre MP / Curitiba MPs → São Paulo MP.

to avoid that, it is required a large sample size to decrease the effect of the sample variance on such tests (Slutsky's theorem). When the normality assumption does not hold, non-parametric statistics should be considered. We propose the use of 3 non-parametric tests: the Kruskal-Wallis ANOVA, the Wilcoxon Signed-Rank test, and the Kolmogorov-Smirnov test. Such non-parametric tests can often have better statistical power in non-normal data.

IV. Evaluation

The objective of the evaluation described in this section is to assess the feasibility and effectiveness of a measurement correlation service considering different statistical tests. First, we depict the employed experimental setup. After that, we describe the experiments performed using these dataset. Finally, a discussion about these experiments is presented.

A. Dataset

The set of traces used to perform the analysis of our proposal are logs produced by the OWAMP *powstream* tool, which measures latency using a continuous stream of one-way tests between two hosts. The logs were collected from the perfSONAR deployment in the Monitoring Service of the IPÊ network (MonIPÊ). Operated by the Brazilian NREN (RNP). The IPÊ network is composed of 27 Points of Presence (PoPs), one in each Brazilian state. Figure 3 shows a graph illustrating the logical topology of the IPÊ network. Each PoP is represented by a square and identified by the name of its host city. The connections are represented by the edges (data transfer rates are described using colors). Besides the represented edges, smaller links connect up to 1200 *campi* and other university locations (e.g., research facilities and hospitals). The IPÊ network is connected to other NRENs, such as the Latin American Advanced Networks Cooperation (RedCLARA)[1] and the american Internet2[2].

The *powstream* data were exported as compressed dumps from the perfSONAR Measurement Archive (MA). These dumps contain measurements results from October, 2014. The *powstream* logs are stored as a structured database in 3 tables: *map_events_powstream*, *powstream*, and *powstream_value_buckets*. The *map_events_powstream* table has measurement metadata (e.g., source/destination IP addresses). The *powstream* table stores general measurements properties metrics (e.g., loss, minimum and maximum delay, and duplicates). The *powstream_value_buckets* table keeps the individual measurements results, which come from the perfSONAR Measurement Points (MPs). The product of the 3 tables (i.e., a SQL INNER JOIN) has approximately 700.000 lines per day. We employ an open data approach, thus the pre-processed dataset is shared through *Harvard Dataverse* [8].

In order to perform the evaluation of the proposed service, we pre-processed the logs through several steps according to the operation of the Transformation Service (TrS) for Measurement Correlation (MC). The first step was to sort the data by time and storing it in a raw CSV file that has only the delay information that we will be using. The second step was to use a moving average for the average delay so that we can match measurement results even when they are recorded at slightly different times. We used a width of 30 minutes for the moving average, producing samples at every 5 minutes when data is available. The pre-processing would be executed by another TrS (considering the perfSONAR architecture). After pre-processing, the TrS for MC retrieves matching data corresponding to the measured delays of two paths inside the network, and analyzes them using parametric and non-parametric correlation statistics.

B. Performed Experiments

The first experiment is the evaluation of the TrS MC considering measurements results from 2 perfSONAR MPs, Porto Alegre and Curitiba, towards another perfSONAR MP, São Paulo (all MPs are represented in Figure 3). We selected these MPs due to regional proximity (i.e., as a convenience sampling). Then, these delay data is used to evaluate the TrS MC using different statistical tests.

We proposed in Section III to use correlation coefficients and distribution tests to perform measurement correlation. Considering such coefficients, the Pearson product-moment correlation coefficient is used as the parametric correlation test and the Spearman rank correlation coefficient and the Kendall rank correlation coefficient are used as the non-parametric correlation tests. The results presented in the Figure 4 consider the correlation of delay data from Porto Alegre MP and Curitiba MP towards São Paulo MP.

[1]RedCLARA - http://www.redclara.net/
[2]Internet2 - http://www.internet2.edu/

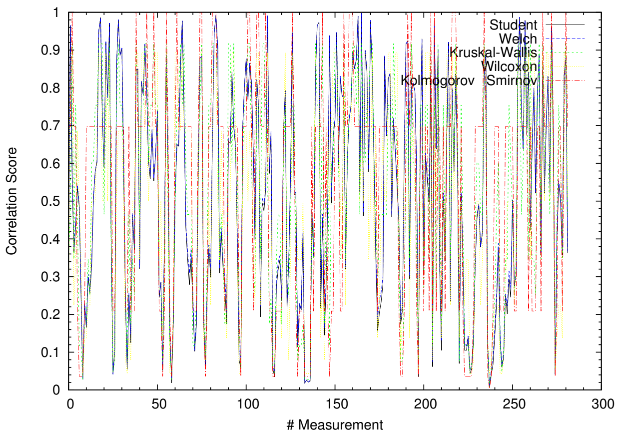

Fig. 5. Distribution Tests - Porto Alegre MP / Curitiba MPs → São Paulo MP.

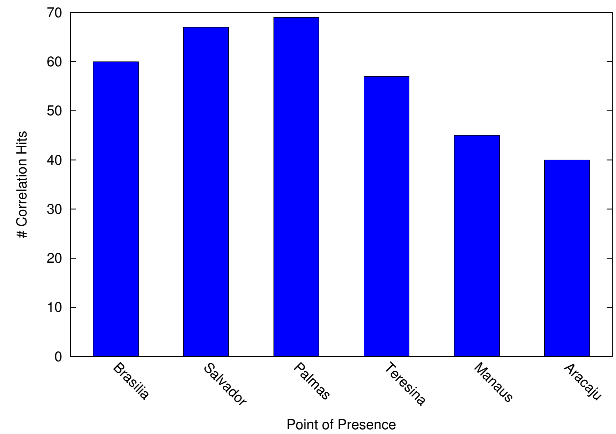

Fig. 6. Summarized Data Analysis - Pearson @ 0.5.

The different correlation coefficients depicted on the curves of Figure 4 present a high concordance. In this context, it is possible to say that any of the coefficient can be used in the TrS MC for the employed dataset. The Pearson product-moment correlation coefficient produce better results when normality holds for sample data, thus, non-parametric tests can be employed in more datasets. On the other hand, since the rank correlation needs the definition of the ranks before the actual calculation of the correlation coefficient, this leads to more costly computations.

The measurement correlation can be performed using distribution tests. Such tests are used to verify whether samples originate from the same distribution. In the present evaluation, we employed the tests described in Section III: the Student's t-test and Welch's t-test, both parametric; and 3 non-parametric tests, the Kruskal-Wallis ANOVA, the Wilcoxon Signed-Rank test, and the Kolmogorov-Smirnov test. The results presented in the Figure 5 consider the correlation of delay data from Porto Alegre MP and Curitiba MP towards São Paulo MP.

The distribution tests depicted on the Figure 5 present a moderate concordance considering the employed dataset. It is possible to say that the parametric t-tests (Student's and Welch's), the non-parametric Kruskal-Wallis ANOVA, and the Wilcoxon Signed-Rank test had similar results. Thus, they can be used in a similar manner in the TrS MC (considering the employed delay data). Regarding the Kolmogorov-Smirnov test, the results diverge from the other tests.

A summarization was performed considering results from an entire day and using a minimum correlation threshold. We chose the Pearson product-moment correlation coefficient as the statistical test since it is popular among correlation tests. Besides that, a threshold of 0.5 to summarize the measurement correlation scores. This threshold was defined considering a related work [4]. When a score is higher than the chosen threshold, we call this a *hit* because it highlights a condition in which the measurement correlation can be used to enable cooperative active measurement tasks. Figure 6 present a summarized data analysis considering the Porto Alegre MP and Curitiba MPs towards different MPs.

The first experiment focuses on the paths departing from Porto Alegre MP and Curitiba MP exclusively. Analyzing all data from Monipê means measuring the correlation between every two paths — and each path is comprised of a source and a destination node. That is, if p is the number of paths, n is the number of network nodes and c is the number of correlations to calculate:

$$\begin{aligned} p &= n \times (n-1) \\ c &= p \times (p-1) \\ &= n^4 - 2n^2 + n \end{aligned} \tag{2}$$

When analyzing the whole network, a further challenge is presenting the information. Figure 7 shows 3 heat maps of correlations between every pair of paths measured in the IPÊ Network: using Pearson (normality assumed), Spearman and Kendall's methods (both non-parametric). Although the figure shows the overview of the amount of correlation present in the network, it fails to provide further insight regarding the relationship between paths without requiring some form of manipulation (axis ordering, for instance). One can also observe that the same shapes appear in all 3 methods, even though the correlation amplitude is quite different.

C. Discussion

The amount of available information about the measurements is essential to define which statistics should be used, parametric or non-parametric. In this context, non-parametric tests are important statistics working tools. Since they consider fewer assumptions about data, they may be applied in more situations, specially when little is known about the dataset properties. In addition, non-parametric tests are usually more robust and simple than parametric ones. On the other hand, when a parametric test is more appropriate (e.g., considering a normal data distribution), non-parametric tests have less power (i.e., they are less accurate).

The required computational resources represent a difference among the employed correlation functions. As widely known, parametric tests are simpler and faster. For example: while Pearson's coefficient calculation has $O(n)$ computational complexity, Spearman's rank correlation coefficient ends up with $O(n \log n)$. This must be taken into account for the complexity of the measurement correlation as whole. However, preliminary tests would likely be required to verify the normality assumptions.

The resource consumption required by the strategy of continuously monitoring and analysing the whole network continuously is prohibitive. On the other hand, whole-network analysis for a short period of time might be required in

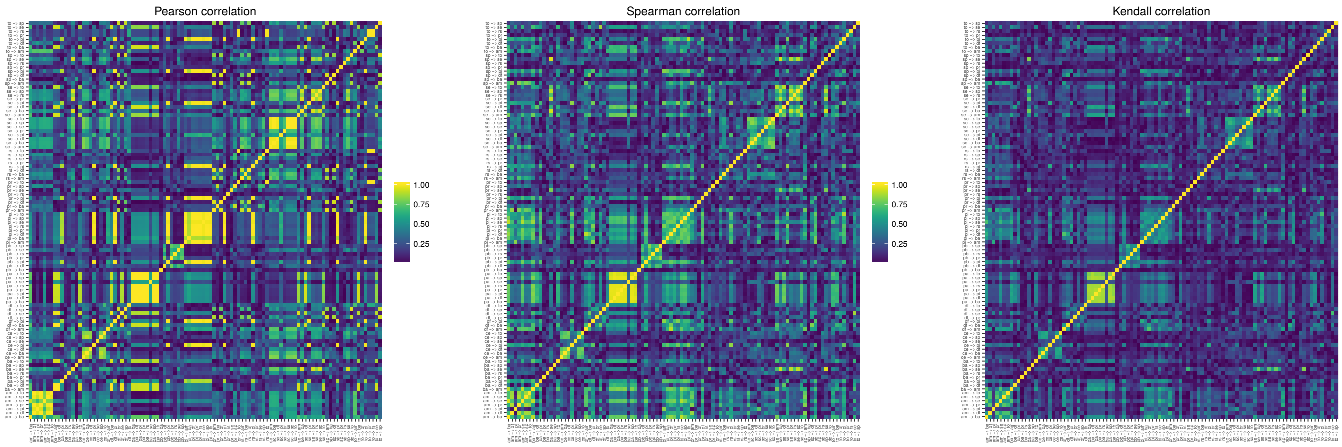

(a) Pearson Correlation (b) Spearman Correlation (c) Kendall Correlation

Fig. 7. Heat maps of measurement correlation in the IPÊ Network.

order to bootstrap the correlations, and allow them to be efficiently monitored continuously by the TrS MC. Nevertheless, it is possible to impose limits on resource usage through the parameterization of algorithms, which enables the measurement correlation process to adapt itself to different network conditions. Finally, intelligent applications can use the resources of measurement federation to work around problems considering the collaboration of nodes.

V. Concluding Remarks

The present work focuses on the utilization of statistical tests to enable measurement correlation in perfSONAR federations as a Transformation Service (Transformation Service for Measurement Correlation - TrS MC). We propose the use of correlation coefficients and distribution tests, both parametric and non-parametric, to implement such service. Since the proposed solution fits into perfSONAR toolkit, it can enjoy the functions performed by other services of such toolkit, e.g., lookup features. The integrated use of the TrS MC in a perfSONAR federation enables the cooperation of different nodes in collaborating organizations for improving monitoring tasks.

The evaluation of the proposed service was performed using an active measurement dataset from the perfSONAR deployment in IPÊ Network (MonIPÊ). The dataset was processed using different correlation coefficient and distribution tests, small scale analysis (PoP to PoP), and large scale analysis (the resulting mesh network which was inferred from the active measurement results). The results show that it is feasible to use the selected statistical tests to correlate measurements in a federation. Besides that, when analyzing the IPÊ network as a whole, it is possible to identify hot spots for measurement correlation. These correlation hot spots can be used as initial screening points for the cooperation of measurement mechanisms inside the federations.

Although the proposed solution shows good results in the performed evaluation, it is necessary to investigate different kind of topologies and conditions. Besides that, our ongoing work is to develop a prototype to integrate the TrS MC into perfSONAR to control the activation of OWAMP sessions in a scalable way. Such prototype can be connected to the perfSONAR Resource Protector service in order to recommend which measurements can be performed through cooperation, due to resource consumption. Finally, measurement correlation could be applied to other active measurement tools, such as the *Bandwidth Test Controller* (BWCTL).

References

[1] B. Claise, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information," RFC 5101 (Standard), Internet Engineering Task Force, January 2008.

[2] S. Shalunov, B. Teitelbaum, A. Karp, J. Boote, and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)," RFC 4656 (Proposed Standard), Internet Engineering Task Force, September 2006.

[3] A. Hanemann, J. W. Boote, E. L. Boyd, J. Durand, L. Kudarimoti, R. Łapacz, D. M. Swany, S. Trocha, and J. Zurawski, "Perfsonar: A service oriented architecture for multi-domain network monitoring," in *Service-Oriented Computing-ICSOC 2005*. Springer, 2005, pp. 241–254.

[4] J. C. Nobre, L. Z. Granville, A. Clemm, and A. G. Prieto, "Decentralized Detection of SLA Violations using P2P Technology," in *Proceedings of the 8th International Conference on Network and Service Management (CNSM 2012)*, October 2012.

[5] J. C. Nobre, L. L. Penz, and L. Z. Granville, "Measurement correlation for improving cooperation in measurement federations," in *Integrated Network and Service Management (IM), 2017 IFIP/IEEE Symposium on*. IEEE, 2017, pp. 584–587.

[6] A. Abdelkefi, Y. Jiang, B. E. Helvik, G. Biczók, and A. Calu, "Assessing the service quality of an internet path through end-to-end measurement," *Computer Networks*, vol. 70, pp. 30–44, 2014.

[7] Y. Zhang, S. Debroy, and P. Calyam, "Network-wide anomaly event detection and diagnosis with perfsonar," *IEEE Transactions on Network and Service Management*, vol. 13, no. 3, pp. 666–680, 2016.

[8] J. C. Nobre, "Owamp powstream tool traces from the perfsonar deployment in the monitoring service of the ip network (monip)," 2016. [Online]. Available: http://dx.doi.org/10.7910/DVN/MSZVYS