Measurement Correlation for Improving Cooperation in Measurement Federations

Jéferson Campos Nobre* Leandro Lisboa Penz†, Lisandro Zambenedetti Granville†,
* University of Vale do Rio dos Sinos, São Leopoldo – Brazil
† Federal University of Rio Grande do Sul, Porto Alegre – Brazil
Email: jcnobre@unisinos.br, lpenz@lpenz.org, 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. In order to ensure that these problems do not occur, service levels need to be monitored. One of the main techniques for this monitoring involves the utilization of active measurement mechanisms. However, such mechanisms are expensive in terms of resources consumption due to the activation of measurement sessions. Thus, such sessions usually can cover only a fraction of what could be measured, which can lead to service level problems being missed. Measurement federations can help network administrators in different tasks, such as the control of 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 which provides measurement correlation. This service can be 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 provides valuable insights regarding measurement correlation from federated measurement points and can be used for the design of better application for the control of active measurement sessions.**

I. Introduction

The capacity and accessibility of networking infrastructures has increased in the last years. Besides that, the diversity of applications and service which operate in such infrastructures has increased importantly. Unfortunately, several times such progress is not directly translated into improved performance for all applications and users. Performance issues can vary from traffic congestion to packet loss caused by intermediate link failures. Such issues lead to financial losses for service providers and customers. Service level monitoring tools must be capable of detecting performance issues in an efficient manner. Measurement mechanisms are one of the most used tools employed for such monitoring.

Service level monitoring is usually performed using either active or passive mechanisms. The network conditions are monitored in a non-intrusive way in passive mechanisms (e.g., IP Flow Information EXport - IPFIX [1]) since monitoring traffic is not created by the monitoring process itself. This is usually done through the observation of network flows inside network devices. Active measurements (e.g., One-Way Active Measurement Protocol - OWAMP [2] and Cisco Service Level Assurance Protocol -IPSLA [3]), on the other hand, are intrusive because they inject monitoring traffic on the network infrastructure to deliver performance metrics. Active measurements can detect end-to-end performance issues in a fine grained fashion. Thus, service level monitoring is usually performed through active mechanisms.

The activation of active measurement sessions is expensive in terms of the consumed resources (e.g., CPU cycles and memory footprint). Besides that, activated measurement sessions also increase the network load due to the injected traffic. At least, one session must be activated to measure a destination, thus the amount of consumed resources is a function of the number of measured destinations. Thus, with more possible paths the larger will be the resources needed to deploy such sessions. The best practice on the operation of active measurement mechanisms is to define the measurement sessions considering the expertise and knowledge of the network administrator. However, current SLA monitoring solutions consume a large amount of computational and human resources, which are not always available to the service provider.

Measurement federations could help network operators to troubleshoot perceived abnormalities as well as improve network middleware regarding service level issues. 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 such mechanisms in important aspects, such as scalability and efficiency. For example, the current best practice regarding active measurements mechanisms 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. 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. In this context, novel approaches are needed to improve the operation of active measurement mechanisms in measurement federations.

Measurement correlation has been used in order to enable efficient control loops for measurement sessions activation [4]. Despite the fact that measurement federations are com-

Jéferson C. Nobre is a former PhD candidate at the Federal University of Rio Grande do Sul, Brazil. He conducted this work during the Monitoring Program (*Programa de Monitoramento*) 2014, funded by the Brazilian National Research Network (*Rede Nacional de Pesquisa*).

posed by programmable nodes, control loops are not used for measurement sessions activation. For example, these loops were applied to improve the detection of SLA violations through an embedded P2P management overlay. In this paper, we present a solution to enable the measurement correlation on federations. This solution enables the coordination of the activation of measurement sessions in order to save resources and decrease the effort required from network administrators.

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

II. BACKGROUND

Active measurement mechanisms are a valuable tool for network administrators. The main goal of such mechanisms is to provide diagnostics which are close to the same network traffic produced by real applications. Some additional design goals considered for the development of active measurement mechanisms are being hard to detect and manipulate, present security features, a logical separation of control and test functionality, and the support for small test packets [2]. Since active measurement results are arguably one of the best ways to determine if a given application will work in a specific network infrastructure, active mechanisms are usually found on network measurement toolkits. Such toolkits usually build measurement federations.

In this section, we first cover one of the most prominent active measurement mechanisms, the One-way Active Measurement Protocol (OWAMP). After that, some widely-known measurement federations are presented.

A. One-way Active Measurement Protocol

One-way measurement protocols are the most straightforward way to detect network issues which happen in one direction of a given network flow. 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 to fetch their results. OWAMP-Test uses UDP to exchange test packets between two measurement nodes. Together, these protocols enable the standardized execution of one-way measurements.

Measurement execution in OWAMP is performed through several different roles. Such roles are logically separated to allow for broad flexibility in use. The first part of the Figure 1 (*a*) shows the logical model used on OWAMP (unlabeled links in this figure means that the respective protocols are unspecified [2]). Session-Sender is the sending endpoint and Session-Receiver is the receiving endpoint of an OWAMP-Test session; Server is an end system that manages one or more OWAMP-Test sessions; Control-Client is an end system that initiates and finishes OWAMP-Test sessions; and Fetch-Client is an end system that initiates requests to fetch the results of completed OWAMP-Test sessions. The different logical roles can be played on different hosts (*a*), but some of these roles can be aggregated in two hosts, as shown in Figure 1 (*b*).

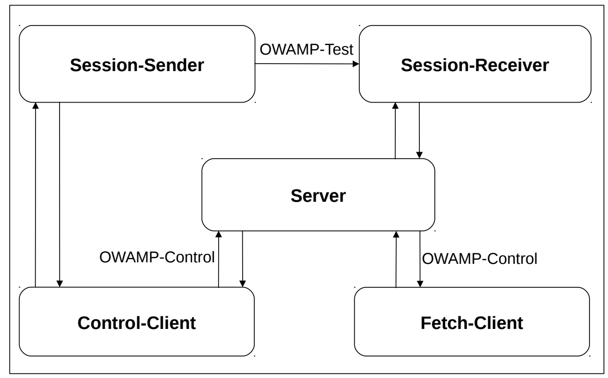

(a) Logical Roles Played by Different Hosts

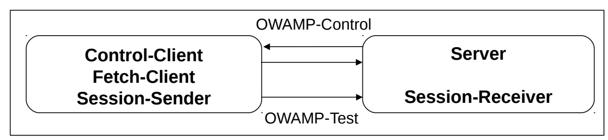

(b) Logical Roles Aggregated in Two Hosts

Fig. 1. OWAMP Logical Models.

The OWAMP tool[1] is an implementation of OWAMP, developed by Internet2 Consortium. This tool is designed to deliver one-way latency, packet loss, packet duplication, and jitter on an end-to-end paths [2]. Small UDP packets are used for the production of measurements (OWAMP-Test) and TCP is employed for test negotiation (OWAMP-Control).

The OWAMP tool comprises several applications for the execution of active measurement tasks, such as *owping* and *owstats*. One of these applications is *powstream*[2], 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. The *powstream* produces data in two formats: raw owamp data files and summary statistics. In addition, powstream data files have the same binary format saved from *owping* and can be parsed using *owstats*.

[1]One-Way Ping (OWAMP) - http://software.internet2.edu/owamp/
[2]powstream - http://software.internet2.edu/owamp/powstream.man.html/

B. Measurement Federations

Sophisticated tools are necessary to monitor large and complex networks since it is required a great deal of effort and expertise from the network administrators. Network measurement toolkits can encompass monitoring and diagnosing tools in order to aid these administrators in their daily tasks. Besides that, such toolkits provides an interface for measurement scheduling, storage of data in uniform formats, and scalable methods to retrieve data. Management stations running such toolkits can be defined as measurement federations.

Measurement federations can aid network administrator teams and increase productivity when utilizing computational and human resources. The focus of these federations is to provide an end-to-end measurement infrastructure capable of network performance monitoring and troubleshooting. Some examples of such federations are the *PERFormance Service Oriented Network monitoring ARchitecture (perfSONAR)*[3], the *SamKnows*[4], the *Grenouille Project*[5], and the *Measurement Lab (M-Lab)*[6]. Some of these federations have thousands of instances deployed world wide, many of which support open data regarding measurement results.

Several measurement federation toolkits are composed of a middleware designed as a Service Oriented Architecture (SOA) to provide federated measurements. In the corresponding federations, services federate with each other, usually locally and globally. In this context, the global federated infrastructure helps to identify and isolate problems in large networked environments. Besides that, services can be designed to automatically discover the presence of other nodes (which also present the same services). The current strongest barrier to the adoption of measurement federations has been interoperability of the toolkits.

There is an inherent human and computational cost related to the deployment of active measurement mechanisms, such as OWAMP, and their continuously operation, i.e., the management of measurement sessions. This cost is strongly related with the size and complexity of network infrastructures. Even using measurement federations, such cost decreases 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.

III. Measurement Correlation in Federations

Open one-way active measurement solutions that are simple to implement and deploy are a valuable tool for network administrators. In this context, inter-domain one-way active measurement could become commonplace [2]. Besides that, one-way measurement tools, such as OWAMP, are usually present in measurement federations. Active measurement sessions must be activated to localize and diagnose performance problems in real-time. However, the activation of such sessions is expensive in terms of consumed resources.

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] [5] [6]. Network devices are correlated peers when they have similar measurement results considering specific metrics. The similarity is quantified through destination scores, which are produced through a comparison of measurement results in a sliding window [4].

We propose the use of destination scores in order to provide correlation features in measurement federations. Such correlation can be used to enable better cooperation features on these federations. For example, specific parts of a network may require more attention due to unexpected events. In this context, measurement nodes can co-operate to reduce operating costs, while providing satisfactory service level compliance. Thereat, the statistics used to correlate the measurement results play an important role in defining which measurement results are similar in terms of the properties required in the sample data (e.g., normality).

In this section, we present a solution to deploy measurement correlation. Such correlation could be deployed in the measurement federation nodes. In addition, we describe some some statistical tests which can be used for measurement correlation. Such tests can be performed through an efficient control of measurement sessions activation.

A. Proposed Solution

We instantiate our proposed solution considering a measurement federation which is defined as a SOA. The employed concept is that each service performs a a limited set of functionality. Some examples of services presented in the measurement federation toolkits are the initiation of performance tests, the storing of performance monitoring results, operations on data sets, registration and location of distributed services, and arbitration on the resources consumption. These services are self contained and provide functionality as either a single or federated deployment. In any case, nodes are designed to automatically discover the presence of other nodes (e.g, using a registration and location service).

In order to support measurement correlation, we describe a service for measurement correlation. Such service perform operations on data sets which have potential to be stored and replayed. Thus, it consumes raw and aggregated active measurement results (such as those produced by OWAMP *powstream*) which are collected and stored by corresponding services. Then, statistical functions are employed to compare collected measurements results. The result of a measurement correlation operation is the evaluation of destination scores. The service for measurement correlation is designed with distributed paradigm in mind, thus it is not necessary to run the measurement correlation on the same station which provides the collection and storage of measurement results.

[3]perfSONAR - http://www.perfsonar.net/
[4]SamKnows - http://www.samknows.com/
[5]Grenouille Project - http://www.grenouille.com/
[6]M-Lab - http://www.measurementlab.net/

The service for measurement correlation is designed to run statistics over changing datasets. Such statistics can be used by network operators to troubleshoot abnormalities and to assist intelligent monitoring and diagnosing tools to remediate network problems (e.g., through network reconfiguration). This can ultimately minimize the impact to end users regarding network faults. Deploying the service for measurement correlation can be done as an isolated application, or through a complete network management solution. For example, external visualization tools, such as Cacti[7] and MRTG[8], can be used to provide reports (e.g., graphics and tables).

The service for measurement correlation must comply with measurement federation policies. One of the possible applications is the activation control of measurement sessions. Such control can help the continuous monitoring of one-way metrics. In addition, the service for measurement correlation interacts through standardized exchange patterns (e.g., using a communication protocol). This can enable a cooperative approach over monitoring problematic destinations through an integrated use of remote network analysis functions.

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 distribution has a certain form, thus, it is possible to infer the value of one or more of its parameters. For example, the hypothesis that a given data distribution has either a specified mean or variance (or both) is statistical. On the other hand, non-parametric statistics make no assumptions about the probability distributions of the variables being assessed.

Measurement correlation can be performed using correlation coefficients produced by different tests. Such coefficients are a measure of statistical dependency between 2 variables, varying between -1 (total negative correlation) and +1 (total positive correlation). The Pearson product-moment correlation coefficient is a parametric correlation test which is used as a measure of the linear dependence between two variables. In addition, non-parametric correlation tests, such as the Spearman rank correlation coefficient can be also used. In this context, some works compared the values of Pearson's and Spearman's correlation coefficients on the same sets of data [7].

The distribution tests are used to verify whether samples originate from the same distribution, usually providing a measure which 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). Since t-tests consider statistical parameters, they are classified as parametric tests. For exactness, it is required normality of the samples (specially when the sample size is small). Besides that, t-tests additionally requires that the sample variance follows a scaled χ^2 distribution. When the normality assumption does not hold, non-parametric statistics should be considered. We propose the use of the the Wilcoxon Mann-Whitney test. In this context, some works compared the choice of either t-test or the Wilcoxon Mann-Whitney for specific statistical metrics [8].

IV. Concluding Remarks

Novel approaches are needed to foster wider adoption of measurement federations as well as enhancing their management capabilities. These toolkits can be conceived as a network-aware middleware for monitoring applications. Measurement correlation can be employed to improve the analysis of measurement data through cooperation features.

The present work focuses on the utilization of statistical tests to enable measurement correlation in SOA-based measurement federations as a service. We propose the use of correlation coefficients and distribution tests, both parametric and non-parametric, in order to implement such service. We use the *powstream* tool, which uses OWAMP as the active measurement protocol, as a model for our proposed solution.

Although the proposed solution shows good properties, it is necessary to evaluate its results in simulation experiments considering measurement federation toolkits. Besides that, our ongoing work is to develop a prototype to use the service for measurement correlation in order to control the activation of OWAMP sessions in perfSONAR. In addition, further analysis could define features and settings required for this prototype, as well as a projection of the resource savings that could be obtained. Finally, measurement correlation could be applied to other active measurement tools, such as the *Bandwidth Test Controller* (BWCTL)[9].

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] M. S. Chiba, A. Clemm, S. Medley, J. Salowey, S. Thombare, and Y. E., "Cisco service-level assurance protocol," RFC 6812 (Informational), Internet Engineering Task Force, January 2013.

[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] ——, "Coordination in P2P Management Overlays to Improve Decentralized Detection of SLA Violations," in *Proceedings of the IEEE International Conference on Communications (ICC 2013)*, June 2013.

[6] ——, "On the Use of Traffic Information to Improve the Coordinated P2P Detection of SLA Violations," in *Proceedings of the 28th IEEE International Conference on Advanced Information Networking and Applications (AINA-2014)*, May 2014.

[7] J. Hauke and T. Kossowski, "Comparison of values of pearson's and spearman's correlation coefficients on the same sets of data," *Quaestiones geographicae*, vol. 30, no. 2, pp. 87–93, 2011.

[8] S. S. Sawilowsky, "Misconceptions leading to choosing the t test over the wilcoxon mann-whitney test for shift in location parameter," *Journal of Modern Applied Statistical Methods*, vol. 4, no. 2, p. 26, 2005.

[7]Cacti - http://www.cacti.net/

[8]The Multi Router Traffic Grapher (MRTG) - http://oss.oetiker.ch/mrtg/

[9]Bandwidth Test Controller (BWCTL) - http://software.internet2.edu/bwctl/