# On the Composition of Performance Metrics in Multi-Domain Networks

*Raphael A. Dourado, Universidade Federal de Pernambuco*

*Leobino N. Sampaio, Universidade Federal da Bahia*

*José A. Suruagy Monteiro, Universidade Federal de Pernambuco*

## ABSTRACT

Despite the advances achieved by service oriented architectures adoption, data originated from measurements are still used in an isolated manner. In multi-domain network scenarios, performance information needs to be processed and related in such a way as to present more relevant results through a metric composition process. Based on recent IETF documents, this article describes a performance metric composition process implemented via software, through a library that can be integrated into diverse visualization and monitoring applications. Experiments using ESnet and Géant backbones performance data demonstrated the effectiveness of the model, the library, and IETF's proposed approach.

## INTRODUCTION

The need to diagnose performance problems in end-to-end communication involving multiple domains led to the deployment of performance monitoring infrastructures, especially in research and academic networks [1]. perfSONAR [2] is an example of such an infrastructure. It is a service oriented architecture able to perform multi-domain measurements and deliver them in a federated environment. The measurements are collected through measurement points running continuous or periodic active tests among them, and made publicly available through measurement archives via Web services interfaces, making it easier to identify and isolate performance related problems even in remote networks. These measurement points can obtain results for metrics such as one-way delay (the amount of time a packet spends in travelling across the network from node A to another node, B), delay variation (the difference in one-way delay between selected packets in a flow), and losses, as well as available bandwidth.

Although these link or intra-domain performance measurements are useful, often it is also necessary to use performance data for the complete end-to-end, usually multi-domain, path. In the latter case, running multi-domain tests may be undesirable due to the extra traffic load, or even impossible due to incompatible measurement infrastructures or policies [3]. Therefore, it may be desirable to reuse the available data in order to produce results equivalent to the ones that would be obtained by running end-to-end tests. Such a procedure is known as spatial metric composition [4, 5].

Even though it is extremely beneficial to both measurement infrastructures and their users, it is noticeable that metric composition is still being underutilized in currently available monitoring and visualization software [1, 6, 7]. These tools, even when accessing a large variety of data from measurement services, do not explore all the possibilities offered by the composition techniques. Besides, most of the work presented in the literature so far is more frequently related to formalizing composition concepts rather than proposing an implementation model [8, 9].

The fact is that handling the measurement data in an isolated way prevents the development of environments that could facilitate the network behavior comprehension and the identification of problems in a given path. Besides promoting a complementary network analysis, metric composition also offers other advantages. It reduces the network overload caused by measurement campaigns and data storage. This happens because usually it is possible to achieve very similar results from previously stored data. Furthermore, metric composition can be useful when trying to obtain a measurement that would be impossible to achieve via an end-to-end test, due to factors such as measurement protocol incompatibilities. Motivated by these advantages, the Internet Engineering Task Force's (IETF's) IP Performance Metrics Working Group (IPPM-WG) had directed its attention toward network

 0163-6804/13/$25.00 © 2013 IEEE 

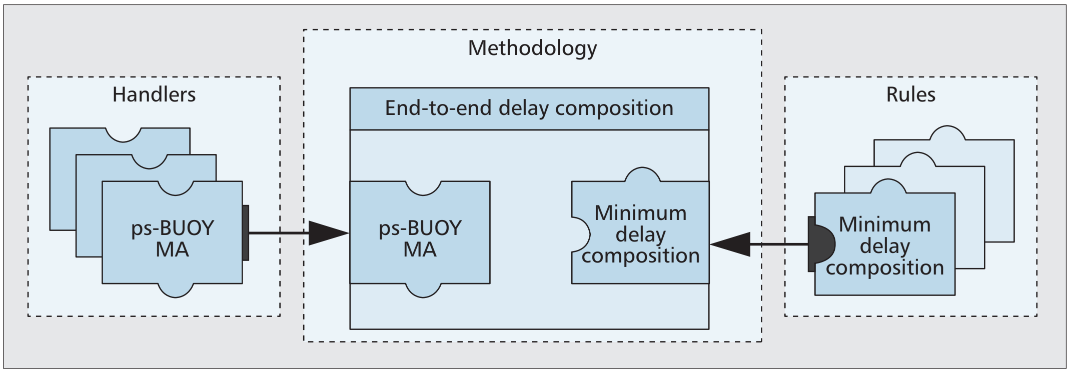


**Figure 1.** *The performance metric composition model.*

*The composition library presented in this article is an implementation of a general model that aims to represent, in a generic and flexible way, the performance metric composition classes in software libraries in view of adapting them to the current monitoring architectures.*

performance metrics composition since the publication of RFC 2330 [10]. Recently, this subject was detailed in RFC 5835 [3], which proposes a framework with general specifications on how to compose metrics, and in RFC 6049 [4], which implements the definitions specified in this framework.

These specifications, associated with the aforementioned needs, motivated this work. While IPPM proposals present general procedures for metric composition, this article introduces the composition of performance metrics through a generic and flexible model capable of encapsulating the metric composition techniques in reusable components. The model allows the inclusion of new composition procedures and techniques whenever they are proposed without requiring major changes in the already developed applications. Furthermore, by using a component oriented approach, monitoring and visualization applications can use this model to perform metric composition in a simplified way.

The remainder of this article is organized as follows. We describe IPPM's definitions related to network performance metric composition. We show the development of the composition metric library and the proposed model, while we describe the results obtained by using the library with real network performance data. Finally, we summarize the work done and present suggestions for future work.

## Network Performance Metric Composition

The composition model described in the next section is based on RFC 5835 [3], which proposes a detailed framework for composing and aggregating network performance metrics defined in RFC 2330 [10].

The main performance metric composition classes proposed in RFC 5835 are temporal aggregation, spatial aggregation, and spatial composition. Temporal aggregation is a class of composition that combines the same type of metric, in which data is obtained in the same scope in different time window intervals. According to RFC 5835, such an approach aims to mitigate the amount of data stored in measurement data archives, thus facilitating the visualization of network characteristics trends. Spatial aggregation also regards the combination of metrics of the same type. However, it performs this activity in different scopes in order to make the performance analysis cover the overall performance of a larger network. The combination of metrics in different scopes requires a weighting scheme to enable the identification of the contribution of each input metric. Finally, spatial composition aims to do performance metric composition of the same type but involving different spatial scopes. The main goal of this composition is to give a potential network performance picture according to direct measurements over a sequence of several spatial scopes.

## The Network Performance Composition Library Development

We developed a software library that implements metric compositions benefiting from the network performance services available in the perfSONAR architecture [2]. The library follows the composition model concepts presented in [11]. The following subsections briefly describe the model and the library components used in the experiments.

### The Performance Metric Composition Model

The composition library presented in this article is an implementation of a general model that aims to represent, in a generic and flexible way, the performance metric composition classes in software libraries, in view of adapting them to the current monitoring architectures. The model is based on a three-step process. The first step consists of defining the way the performance data is obtained. That is, it is dedicated to searching the performance data sources and measurement services. The second step comprises data gathering, available through one or more measurement services. It requires knowledge of the measurement interfaces used to make such data available to software applications. The third step regards composition. It processes the measurement data according to a given rule previously defined to generate summarized results.

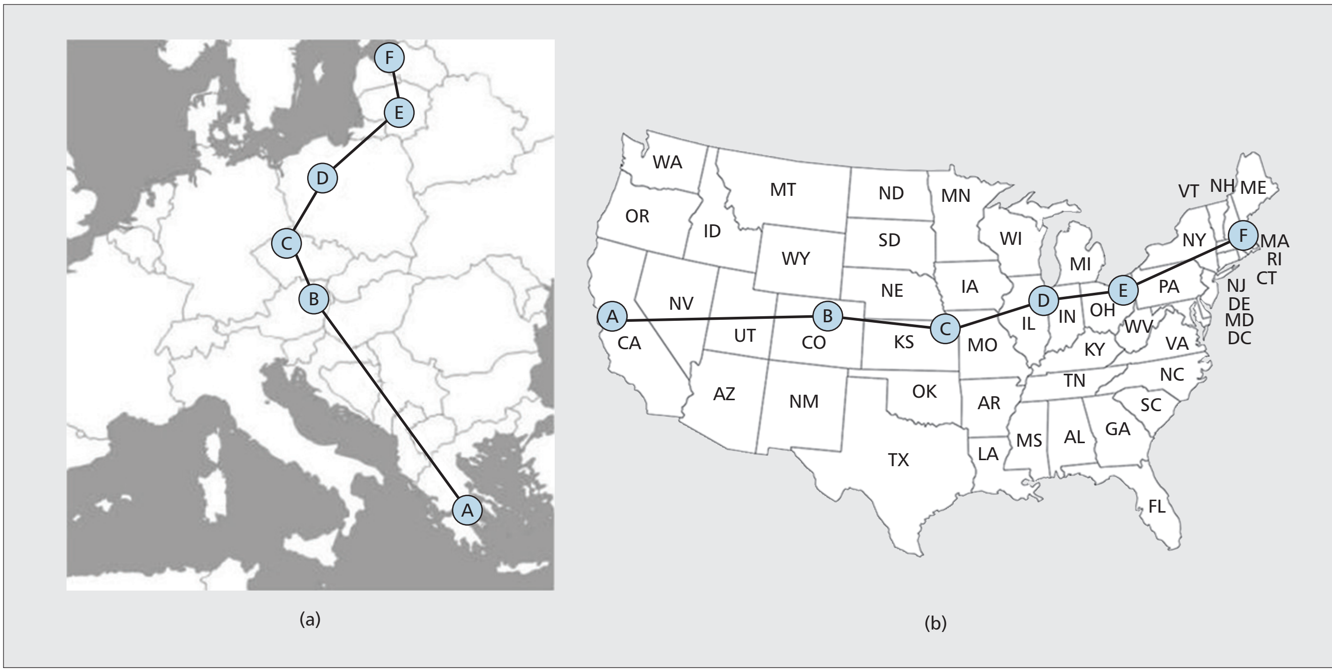


**Figure 2.** *Route used at a) Géant; b) ESnet networks.*

The model depicted in Fig. 1 abstracts such steps through three software components: a composition methodology, one or more handlers, and a composition rule. These components are combined to represent the diverse types of performance compositions proposed in RFC 5835. The following subsections present the development details of the composition library, which followed such a model.

### The Composition Library

The composition library was developed in Java to furnish monitoring applications with all the benefits of the IPPM proposed metric compositions in a simplified way. This library was tested using performance data available at Energy Sciences Network (ESnet) and Géant (pan European network) through bundles coupled to the ICE tool [6]. ICE is a visualization tool based on the FLAVOR framework [12] and designed to provide network performance data visualization in an integrated, flexible, and reusable way, through software packages called bundles. The experiment results (shown in the next section) enabled us to verify IPPM's composition approach, the composition model, and the developed library.

Besides the proposed composition model, Fig. 1 also presents a use case of the implemented library. It shows the model components used to spatially compose an end-to-end delay. This type of composition requires the use of one or handlers capable of providing the required delay data for each sub-path. Besides that, it also requires a rule, belonging to the spatial composition class, which is capable of composing delay data. In this example, the *pS-BUOY MA* [13] handler and the *Min Delay* composition rule were chosen.

## Experiments and Results

The experiments described in this section aimed to verify the distortion between end-to-end measurements and the results obtained by composing sub-path measurements. In order to accomplish this, we fixed the start and end points of each path to be used in the experiments as well as the route used at the time of the measurement. Thus, we collected data from each section that makes up the full path, within the same time interval of the end-to-end measurement, and then used this data to perform the composition.

Following this procedure, we conducted experiments in two different measurement environments. The first one was Géant's network, which uses a service-oriented architecture based on perfSONAR and maintains measurement points running the HADES system [14], which performs the required tests, collects the results, and stores them in a centralized measurement archive to make these data available.

The second network environment used was ESnet. This network also uses a service-oriented architecture based on perfSONAR and maintains measurement points using the perfSONAR-PS services [15] to manage performance testing. The measurement archives of this network are stored in a distributed fashion, and the delay data made available through the pS-BUOY measurement archive service [13] running on each measurement point.

Based on the environments described above, we implemented the necessary components in the composition library presented earlier. Therefore, we created the handlers HADES MA and pS-BUOY MA to retrieve either the complete path (end-to-end) data as well as the measurement data related to each sub-path. In addition, we implemented the composition rules *minimum delay spatial composition*, *mean delay spatial composition*, and *delay variation spatial composition*, as well as the composition methodologies *delay spatial composition* and *delay variation spatial composition*.

The following subsections present the results obtained using delay composition and delay variation composition, respectively.

### DELAY COMPOSITION

As stated in [4], the delay spatial composition involves the concatenation of measurements collected from each sub-path of the complete path. Therefore, the composed delay can be obtained by summing the corresponding mean or minimum delay observed in each sub-path that composes the desired complete path.

Following this definition, the first experiment used Géant's measurement data from the path between one node in Athens, Greece, to another one in Riga, Latvia. The route used by these nodes between 0 h and 23:59:59 on December 9, 2010 was obtained from HADES's measurement archives, and is illustrated by Fig. 2a. The reported route was Athens (A) — Vienna, Austria (B) — Prague, Czech Republic (C) — Poznan, Poland (D) — Kaunas, Lithuania (E) — Riga (F).

The average delays obtained from both the end-to-end measurements and the composition by using the route described in Fig. 2a are shown in the upper part of Fig. 3. All historical data were fetched from the HADES measurement archive at the University of Erlangen-Nuremberg. Figure 3 shows a comparison chart in which the vertical axis represents the average delay in milliseconds, the horizontal axis represents the sequence number of the sample, the solid lines show the end-to-end measurement data, and the dotted lines represent the results from the composition. The time window used provided a total of 1440 samples taken at 1-min intervals.

The second experiment used ESnet's data from the path between one node in the city of Sunnyvale, California, toward another one located in Boston, Massachusetts. The route that was being used by these nodes between 0 h and 23:59:59 of December 13, 2010 was obtained from the perfSONAR-PS Traceroute Measurement Archive service, and is illustrated by Fig. 2b. The reported route was Sunnyvale (A) — Denver, Colorado (B) — Kansas City, Kansas (C) — Chicago, Illinois (D) — Cleveland, Ohio (E) — Boston (F). The lower part of Fig. 3 shows the composed and end-to-end minimum one-way delay results for experiments conducted in this environment, retrieved from the measurement archives stored at each node. Based on the results presented in this subsection, it is clear that the delay spatial composition is capable of producing trustworthy estimations. It is also noticeable that the results obtained through the minimum delay composition are slightly more accurate than those obtained through the mean delay composition, which is expected since IP networks' minimum delays usually present little variation.

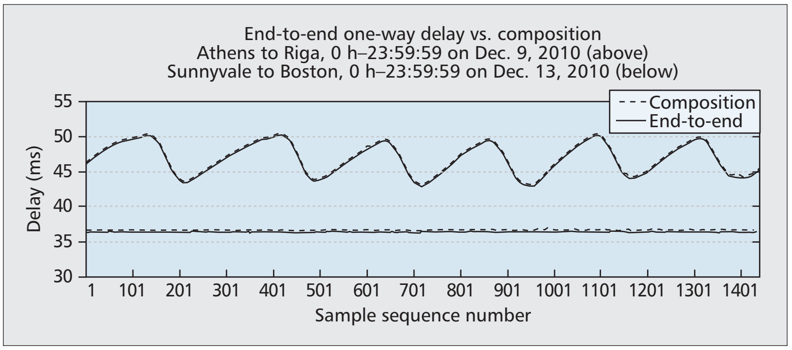


**Figure 3.** *Comparison between the end-to-end delay measurement and composition.*

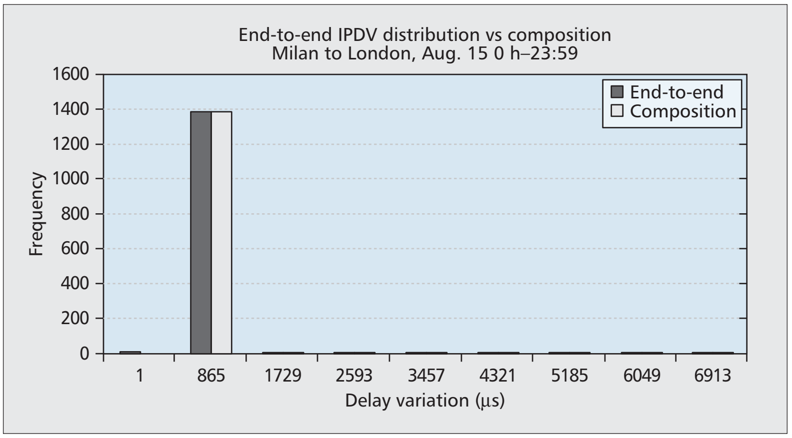


**Figure 4.** *IPDV distribution (Milan – London, August 15, 2011).*

### DELAY VARIATION COMPOSITION

Among all performance metrics, the delay variation shows the lowest accuracy when subjected to composition [5]. This characteristic is a result of the unavailability of detailed information about the delay distributions at each portion of the network. In the face of these problems, alternative solutions have been proposed in RFC 6049, which presents two methods for performing the delay variation (DV) composition, based on preliminary studies reported in the International Telecommunication Union — Telecommunication Standards Sector's (ITU-T's) Recommendation Y.1541.

In this work, we used the method described in [5], which is based on the properties of the normal power approximation. The overall goal is to estimate a quantile $t$ of the end-to-end delay $T$, such that $Pr(T < t) = p$, assuming $p$ to be the 99.9th percentile ($p = 0.999$). Based on this approach, it is possible to estimate the delay variation as the difference between the 99.9th quantile and the minimum delay observed across the evaluation interval.

Following these definitions, we conducted several experiments in order to verify the accuracy of this type of composition. The first one used Géant's data from nodes between Milan and London, from 0 h to 23:59:59 on August 15, 2011. The route that was being used was Milan — Geneva — Paris — London. The results are shown in Fig. 4 through a histogram that compares the distribution of IPDV values obtained

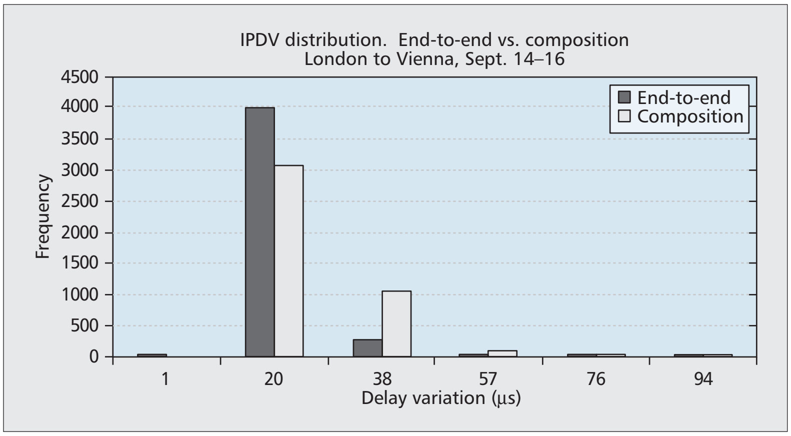


**Figure 5.** *IPDV distribution (London — Vienna, Sept. 14–16, 2011).*

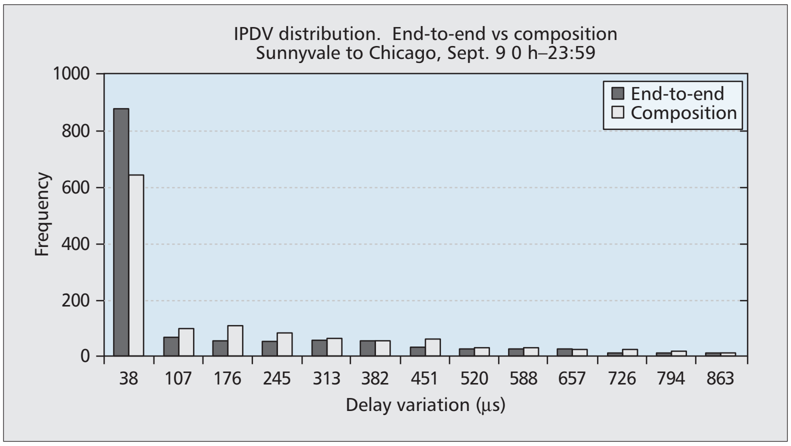


**Figure 6.** *IPDV distribution (Sunnyvale — Chicago, Sept. 9, 2011).*

via the composition and those retrieved from end-to-end measurements.

In order to evaluate the method in a longer time period, we applied the same composition technique using data from another section of the same network, with a 48-h time frame (between 0 h of September 14 and 23:59:59 of September 16, 2011). The path used was London — Amsterdam — Frankfurt — Prague — Vienna. The comparison between the end-to-end and composed IPDV is shown in Fig. 5. As can be seen, even in a larger interval, the results provided by composition are reasonably close to the ones obtained via end-to-end measurements.

We also conducted IPDV composition experiments using ESnet's data, using the path from Sunnyvale to Chicago, in two different time intervals. The results obtained in the first experiment (between 0 h and 23:59:59 of September 9, 2011) are shown in Fig. 6. We conducted another experiment using data from September 14, and the results were very similar to those of the first one. The route used in both experiments was Sunnyvale — Denver — Kansas — Chicago.

Based on the previous results, one can observe that the distribution of the IPDV produced by the composition remains reasonably close to that presented for the end-to-end measurements. Although this particular type of composition is not as accurate as those presented in the previous section, it is possible to estimate with some certainty the delay variation over the analyzed period.

## Conclusions

This article showed the applicability of network performance metric composition in large Internet communication backbones such as ESnet and Géant. The findings of the experiments demonstrate that the results obtained through composition can assist network analysis through complementary information regarding the performance of a network path involving multiple domains. Such information is not always available when the measurement practitioner needs it, particularly in wide scopes such as the Internet. This work presents a practical use case of performance metric composition that shows its feasibility in real scenarios through the adoption of perfSONAR measurement services. As future work, it is intended to conduct experiments with the packet loss composition — which was not done yet due the difficulty of finding relevant loss rates in research high-speed networks — and compositions across multiple domains.

### Acknowledgment

The authors would like to thank the Brazilian National Research and Education Network (RNP), the Brazilian National Council for Scientific and Technological Development (CNPq), and Bahia's State Research Foundation (FAPESB) for all their support during the development of this research. The authors would also like to thank ESnet and Géant for making available the measurement data for their networks, especially Sowmya Balasubramanian from ESnet, and Roland Karch and Stephan Kraft from Géant.

## References

[1] L. Sampaio *et al.*, "Implementing and Deploying Network Monitoring Service Oriented Architectures," *5th Latin American Network Operations and Management Symp.*, Sept. 2007.

[2] A. Hanemann *et al.*, "PerfSONAR: A Service Oriented Architecture for Multi-Domain Network Monitoring," *Service-Oriented Computing ICSOC 2005*, LNCS, vol. 3826, 2005, pp. 241–54, http://www.springerlink.com/content/f166672085772143/.

[3] A. Morton and S. V. den Berghe, "RFC 5835: Framework for Metric Composition," Apr. 2010, http://www.rfceditor.org/rfc/rfc5835.txt.

[4] A. Morton and E. Stephan, "Spatial Composition of Metrics," IETF RFC 6049 (Proposed Standard), Jan. 2011, http://www.ietf.org/rfc/rfc6049.txt.

[5] ITU Rec. ITU-T Y.1541, "Network Performance Objectives for IP-Based Services," Feb. 2006.

[6] I. Koga *et al.*, "Ice: A Flexible Network Monitoring Access Environment," *Tool Session, 25th Brazilian Symp. Computer Networks and Distrib. Sys.*, 2007.

[7] A. Hanemann *et al.*, "Complementary Visualization of Perfsonar Network Performance Measurements," *Proc. Int'l. Conf. Internet Surveillance and Protection*, IARIA/IEEE, Aug. 2006.

[8] V. Paxson, "Towards A Framework for Defining Internet Performance Metrics," *Proc. INET '96*, 1996.

[9] A. Hanemann *et al.*, "A Study on Network Performance Metrics and Their Composition," *CampusWide Info. Sys.*, vol. 23, issue 4, 2006, pp. 268–82.

[10] V. Paxson *et al.*, "RFC 2330: Framework for IP performance Metrics," May 1998, Status: Informational; ftp://ftp.internic.net/rfc/rfc2330.txt, ftp://ftp.math.utah.edu/pub/rfc/rfc2330.txt.

[11] R. A. Dourado *et al.*, "A Model for Network Performance Metric Composition (in Portuguese)," *Proc. 29th Brazilian Symp. Computer Networks and Distrib. Sys.*, Campo Grande, MS, 2011.
[12] I. Koga, L. Sampaio, and J. A. S. Monteiro, "Flavor: A Dynamic and Open Framework for the Development of Network Measurement Access and Visualization Tools," *25th Brazilian Symp. Computer Networks and Distrib. Sys.*, 2007.
[13] Internet2, "Perfsonar-Buoy Measurement Archive," http://psps.perfsonar.net/psb/, Apr. 2013.
[14] P. Holleczek *et al.*, "Statistical Characteristics of Active IP One Way Delay Measurements," *Networking and Services*, 2006, p. 1.
[15] Internet2, "Network Performance Toolkit," http://psps.perfsonar.net/, Apr. 2013.

## BIOGRAPHIES

RAPHAEL A. DOURADO (rasd2@cin.ufpe.br) received his B.S. in computer science from Universidade Salvador (UNIFACS), Brazil, in 2011. He is currently an M.Sc. student at Universidade Federal de Pernambuco (UFPe), Brazil.

LEOBINO N. SAMPAIO (leobino@ufba.br) received his B.S. and M.Sc. degrees in computer science from UNIFACS, and his Ph.D. degree in computer science from UFPe, in 1997, 2006, and 2011, respectively. Since March 2010, he has been a faculty member at Universidade Federal da Bahia, where he is currently an adjunct professor. His research interests include the areas of network performance measurements, future Internet, and software-defined networks.

JOSÉ A. SURUAGY MONTEIRO (suruagy@cin.ufpe.br) received his B.E. in electrical engineering from UFPe in 1979, his M.Sc. degree in electrical engineering from Universidade de São Paulo, Brazil, in 1982, and his Ph.D. in computer science from the University of California, Los Angeles in 1990. He is currently a professor, researcher, and vice-director at Centro de Informática, UFPe. His research interests include computer network architecture, traffic engineering, network measurements, quality of service, and performance evaluation.