

Quality of Service monitoring adopting correlation among active and passive measurements: The experience from the FP7 mPlane project

Original

Quality of Service monitoring adopting correlation among active and passive measurements: The experience from the FP7 mPlane project / Mellia, Marco; Traverso, Stefano; Tego, E.; Rufini, A.; Valenti, A.; Matera, F.; Scarpino, M.; Kahveci, A. E.; Invernizzi, F.. - ELETTRONICO. - (2016), pp. 176-181. (Intervento presentato al convegno 17th International Telecommunications Network Strategy and Planning Symposium, Networks 2016 tenutosi a Montreal, CA nel September 2016) [10.1109/NETWKS.2016.7751172].

Availability:

This version is available at: 11583/2666871 since: 2017-03-11T14:26:58Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/NETWKS.2016.7751172

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Quality of Service monitoring adopting correlation among active and passive measurements: The experience from the FP7 mPlane project

M. Mellia*, S. Traverso*, E. Tego**, A. Rufini**, A. Valenti**, F. Matera**, M. Scarpino***, A. E Kahveci***, F. Invernizzi****.

* Dip. di Elettronica e Telecom., Politecnico di Torino, Corso Duca Degli Abruzzi 24, 10129, Torino, mellia@tlc.polito.it

** Fondazione Ugo Bordoni, via del Policlinico 147, 00161 Roma, mat@fub.it

*** FASTWEB, Via Caracciolo, 51, 20155 Milano – Italy, ahmeterhan.kahveci@fastweb.it

**** TIM, Core Network and Infrastructure IP Edge Innovation, Via G.Reiss Romoli, 274 10148 Torino, fabrizio.invernizzi@telecomitalia.it

Abstract—In this work we report the final results of the field demonstration of the European FP7 mPlane project concerning the evaluation of the network performance with particular details regarding the access Quality of Service and traffic monitoring. We show how the proposed and experimented architecture of this measurement plane is fundamental for many future internet evolutions concerning user perception, content popularity web quality and network management. Here we point out the results dedicated to service level agreement verification and certification with analysis of the correlation between passive and active measurements in order to understand some causes concerning network and service anomalies.

Keywords—QoS; SLA, GPON; network congestion; TCP; UDP.

I. INTRODUCTION

The distributed nature of internet leads to operation brittleness and difficulty in identify and tracking the root causes of performance and availability issues. Looking at these aspects it can be difficult, or even impossible, to assess why a given problem is occurring, especially at scales that encompass several network entities and devices. Even understanding the cause of the quality degradation has become a daunting task, made even more challenging by the fast and constant deployment of new services and applications. This makes it difficult to take the correct countermeasures when issues arise, thus limiting network management and operational activities. In short, the Internet is often a large, obscure black box, and Network and Application Providers as well as End-Users lack the necessary mechanisms to verify the offered service level and to drill down to the cause of problems therein.

To avoid such internet limitations and to exploit all the network capacities the FP7 MPLANE (*an Intelligent Measurement Plane for Future Network and Application Management*, <http://www.ict-mplane.eu/>) project proposed to introduce a measurement plane alongside the Internet's data and control planes, able to collect the measurements from several probes,

both active and passive, located in different points of the network, with an intelligent reasoner that permits to analyze all the corresponding data, summarizing and illustrating output information in terms of network performance, traffic characteristics, Quality of Service (QoS) and user perception.

During the duration of such a project (2012-2015) several activities and use cases were investigated also with several field trials and details can be found in the deliverables reported in <http://www.ict-mplane.eu/>. In this paper we report our experience on the theme of the QoS and network performance and in particular with problems regarding the protection of Internet users' rights. It has to be pointed out that these themes have a fundamental role for the relationships between Internet Service Provider (ISP) and clients and National-level agencies, charged with protecting such rights (for example AGCOM [1] in Italy, or OFCOM in UK), have already adopted specific solutions in verifying that the promised performance is delivered and Service Level Agreement (SLA) are met, even though some difficulties have been found, especially when targeting high-speed access, e.g. capacity higher than 20Mb/s [2]. Defining what and how to measure in these scenarios, and sorting out why or who is responsible when things fall short of expectations, is complex. Tests must be continuously updated to follow the introduction of new applications, and must go beyond simplistic network-layer indexes like round-trip-time, bulk download rate and packet loss. The solution adopted in this project for SLA verification and certification is based on the use of a novel probe, mSLAcert [3], that permits to evaluate the line capacity by means of UDP tests, compared with TCP ones [4]. Such a probe has been tested in a xDSL-GPON (Gigabit Passive Optical Network) lab [4] and details are reported in the mPlane deliverable D5.6 [5]. Here we demonstrate the reliability of mSLAcert in a wide European network. Furthermore the network performance is analyzed in terms of a correlation between active and passive tests for an ADSL access network

belonging to an ISP, *Fastweb*, partner of this project.

This paper is structured in the following way: after this introduction in Sect. II the mPlane architecture is described, in Sect. III the mSLAcert measurements are reported in the context of a wide European network, comparing TCP and UDP tests; in Sect. IV the analysis about the correlation among passive and active measurements is described for an ADSL field environment. Finally conclusions are drawn in Sect. V.

II. MPLANE ARCHITECTURE

MPLANE consists of a Distributed Measurement Infrastructure to perform active, passive and hybrid measurements at different OSI layers.

A *repository* and *analysis* layer collects, stores, and analyses the collected data via parallel processing and data mining. Finally, an *intelligent reasoner* iteratively drills down into the cause of an evidence, determining the conditions leading to given issues, and supporting the understanding of problem origins. In particular the *measurement layer*, which combines a set of new (software and hardware) programmable mPlane probes with legacy probes adapted to the mPlane measurement layer interface into a common, large-scale, distributed measurement layer, which is concurrently accessible to a wide set of stakeholders (ISPs, Application Providers, Regulators, Researchers, etc.).

One of the main target of mPlane is to assist the user in all his requirements concerning its access to the network, with particular regards to QoS tests and to verification and certification of the SLA between user and ISP [3]. But in order to analyze troubles more related to user perception (Quality of Experience, QoE) [6] network tests related not only the access part, but also to the inner one (metro and core) are necessary, especially to investigate congestion events. For such an aim the

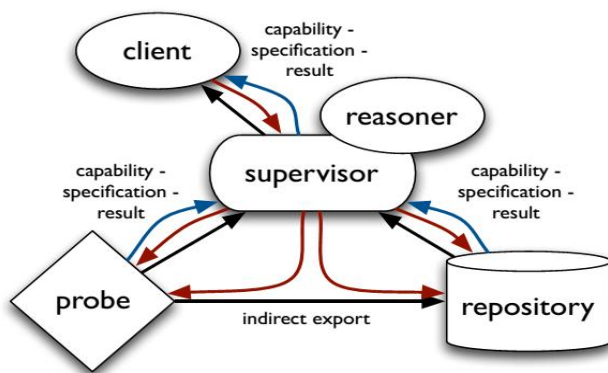


Fig. 1. The mPlane architecture. The Reasoner coordinates the measurements and the analysis performed by probes and repositories, actuating through the Supervisor.

mPlane measurement layer results fundamental since it provides a distributed and ubiquitous network monitoring framework [7] to gather heterogeneous measurements from an assorted number of different vantage points.

Fig. 1 recalls the architecture of the mPlane. The Reasoner coordinates the measurements and the analysis performed by probes and repositories, actuating through the Supervisor. It is responsible for the orchestration of the iterative analysis and the correlation of the results exposed by the analysis modules. Such a reasoning-based system is capable of generating conclusions and triggering further measurements to provide more accurate and detailed insights regarding the supported traffic monitoring and analysis applications. As such, the reasoner offers the necessary adaptability and smartness of the mPlane to find the proper high-level yet accurate explanations to the problems under analysis in the different use cases. As example we can invite the reader to deepen some topics developed in the mPlane use cases as described in <http://www.ict-mplane.eu/public/use-cases>. In particular in the final deliverable D6.3 [8] several results are reported concerning: estimating content and service popularity for network optimization, anomaly detection and root cause analysis in large-scale networks, passive content promotion and curation, active measurements for multimedia content delivery, QoE for web browsing, mobile network performance and SLA verification and certification; this last use case will be treated in the following.

SLA verification and certification

SLA between providers and customers of Internet services regulate the minimum level of service provided in terms of one or more measurable parameters. The verification of an SLA is technically equivalent to the verification of the implicit guarantees of service made by a provider offering Internet service. SLAs are generally tested in terms of some network performance parameters as "bandwidth" (generally expressed in terms of raw throughput). However, with the evolution of the applications, SLA will regard aspects more and more related to user perception.

Currently most of the SLA verifications are based on speed tests that measure TCP throughput. However, since from its introduction, TCP was well known for its limits in terms of bandwidth exploitation [9], but so far this aspect was not in discussion since the available access capacities were the ones mainly offered by the ADSL-ADSL2 lines, with values lower than 20 Mb/s [2]. Conversely in recent years the available bandwidth for users has become so great as to call into question the performance of TCP in networks with high bandwidth-delay product. Recently new TCP versions have been introduced, with the aim to better exploit the bandwidth at disposal, however there still exist several situations in which the bandwidth is

never exploited up to 100% [10]. The measurement of the line capacity for high bandwidth-delay product links is a debate that has lasted years, but with the growth of the optical fiber penetration such a topic becomes very urgent. So, in the framework of mPlane project a novel method was adopted, based on a single UDP session [4], and the reliability of such an approach was experimentally verified in the case of GPON networks. Based on such a method a novel probe, mSLAcert, was implemented and tested in lab and details are reported in mPlane Deliverable D5.6 [5]. It is composed of two components: a server and an agent. The measurement is based on RTT tests and TCP/UDP downloads from server to agent; at end of the download tests the agent sends a report back to the server, reporting the measured parameters. The reasoner can request additional tests for the SLA verification to the supervisor. The supervisor sends the test specifications to the probe, which after completing the test it will send the result back to the supervisor, that forward it to the reasoner. After the result comparison carried out by the reasoner, results are presented to the user as signed PDF document to certify the delay and throughput measured by the agent, in such a way verifying the SLA. In the following we report the field tests carried out for mSLAcert during the final demonstration in Heidelberg.

III. SLA VERIFICATION AND CERTIFICATION IN A LONG SCALE SERVER CLIENT NETWORK

Among the several results carried out in the framework of such a project we point out the final demonstration that was tested in a wide European network environment shown in fig.2, with reference to the SLA verification and certification use case. Four testbeds have been interconnected, one in Heidelberg (hosting the client and server probes), one in Fastweb premises in Milano (hosting the supervisor, and the server probe), one in Telecom Italia testpant in Torino (hosting another server), and one in FUB Laboratories in Roma (hosting another client and server). In particular, such a test verified that the line capacity can be correctly measured also in *high bandwidth x delay* product links by adopting the UDP procedure defined for the mSLAcert probe [3], also by considering very long distance connection between the server and the client, and overall crossing networks of different operators and Countries. The components SLA Server and SLA Agent, were installed on four points, FUB, NEC, Fastweb and Telecom Italia. All the components were integrated with the public Supervisor, located in Fastweb premises in Milano. Furthermore, for these demo purposes, on the connection reserved for this use case in Fastweb, some network impairments were added. These allow us to test the system in adversarial conditions, for instance increasing the delay, and introducing artificial packet loss.

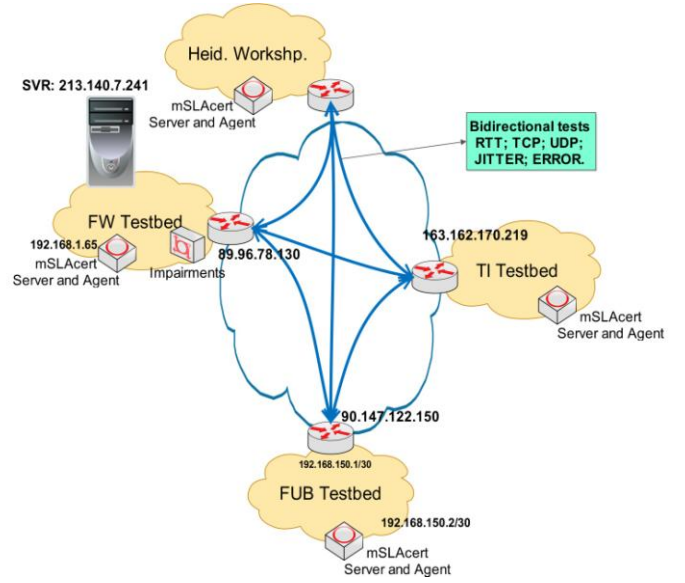


Fig. 2. the mPlane architecture operating in Europe environment.

In fig. 3 we show a typical screenshot from mPlane tool for SLA certification illustrating the throughput time behavior in the case server-client corresponding to Milan-Heidelberg, where 100 Mbps were reserved on both access sides. The figure shows that, by adopting TCP tests, the maximum exploiting throughput was 50 Mbps, while the line capacity was 100 Mbps. This was expected, since a single TCP connection cannot saturate 100 Mbps physical layer capacity in a long distance scenario, where the RTT was more than 30ms.




Fig.3. TCP throughput behavior (Mb/s) in the Heidelberg-Milan (server-client) link.

msla-detail-ip4-15

Parameters			
Name ↑	Value		
destination.ip4	195.37.61.86		
source.ip4	89.96.78.130		
when	2015-11-30 11:16:10.238198 ... 2015-11-30 11:16:18.669595		
<div><div>Values</div><div>Chart</div><div>Add to On-Demand</div></div>			
time	delay.twoway.icmp	mSLA.tcpBandwidth	mSLA.udpCapacity
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	99
2015-11-30 11:...	167000	43	92

Fig. 4. Screenshots from mPlane measurement comparing TCP and UDP throughput tests (Mb/s) in the Heidelberg-Milan (server-client) link.

In fig. 4, for each test, we compare TCP and UDP throughput for the same server-client environment of fig. 3 and we confirm that by adopting UDP test the measured throughput was always close to the line capacity, conversely the TCP showed a value of 43 Mb/s. Recall that the UDP test use a greedy heuristic to find the maximum rate that allows to exchange data without losses on the path. By getting rid of TCP congestion control algorithms, the UDP greedy source is able to effectively test and measure the available bandwidth [3].

IV. CORRELATION AMONG ACTIVE AND PASSIVE TESTS IN AN ADSL ENVIRONMENT

A first investigation on correlation among passive and active measurements was reported in [7], where a clear relationship among bandwidth, retransmitted packets and RTT was found in ADSL tests in the presence of low signal to noise ratio and network congestion. In this paper the correlation was analytically investigated in terms of the Pearson covariance. We adopted the field ADSL tests carried out on the Fastweb network, with active probes located in more than 30 places scattered in the operational country-wide network, during a period of three months. Details on this field trials have already been reported in [7]. In particular, we adopt the Internet QoS Measurement (IQM) probe developed by Fastweb, scattered in the ISP network that are instrumented to periodically download (upload) via FTP some predefined files from (to) a server, logging the application layer throughput. The machine hosting the FTP server runs also Tstat [11], a passive sniffer that extracts transport-layer information about the TCP flows. The IQM probe consists of several components glued together by a modular architecture as illustrated in fig. 5. Each tests lasted at




Fig. 5. Measurement network at Fastweb ISP.

least 10 seconds and the following results refer to Downlink ADSL.

This investigation is based on the covariance approach as defined by means of the correlation Pearson coefficient. Such a correlation expresses the degree of relationship between two variables, x and y , and it is measured by the correlation coefficient " r "


$$r_{xy} = \frac{\sum_{i=1}^n x_i - \bar{x} \quad y_i - \bar{y}}{\sqrt{\sum_{i=1}^n x_i - \bar{x}^2} \sqrt{\sum_{i=1}^n y_i - \bar{y}^2}} \quad (1)$$

$$\{-1 \leq r \leq +1, \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i, \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i\}$$

n - Nr. of elements of dataset $x_1, x_2, \dots, x_i, \dots, x_n, y_1, y_2, \dots, y_i, \dots, y_n$

which ranges from -1 to 1. Where the sign indicates the direction of the relationship. So, the correlation could be positive when the values of two variables are changing with same direction; and could be negative when the values of variables are changing with opposite direction.

For our aims, we decided to assume as variables under observation the mean bandwidth (throughput), B , of a flow, the number of retransmitted segments of the flow due to time out (R_{TO}) and the ones due to fast retransmit (R_{FR}) process. The methodology is based on the following assumption: when we have a network trouble the anomaly is reflected in some distinguished behaviors of the monitored data and will impact more than one variable of the transmission flow. Here such a methodology is applied in the case of network congestion, that strongly influences on the bandwidth, timeout and fast retransmit processes.



In fig. 6 the behavior of a reference case is reported where no particular problem was observed during the monitoring. In such a case, the bandwidth is stable, around 9.5Mbps.

The active measurement is based on 20 tests run every hour, and for every test also the passive monitoring was analyzed in terms of number of retransmitted packets due to time out and fast retransmit.

In fig. 7, we show, the passive monitoring corresponding to the throughput of fig. 6. In the same figure we also include a two raw table reporting the values of the correlation r between B and R_{TO} (up) and B and R_{FR} (below). It has to be pointed out that in the presence of constant results in time, with no correlation between two variables, the r coefficient tends to assume a 0/0 form and therefore in such a case we report a *NaN* value. In fact, in this case, the correlation has the same value, zero, for almost all the period of 24 hours, which means that no relationship exists between B vs R_{TO} and vs R_{FR} . However also for this well stable QoS case we can notice a weak anomaly, and in fact, directly from the correlation, we can notice a negative (-0.9) value, in the red area due to a number of retransmissions caused by time out. Such an anomaly corresponds to the small bandwidth degreasing indicated in the red box in fig. 6.




Fig. 7. Correlation of the passive data for the case of fig. 6.




Fig. 8. Throughput from the active monitoring, case with congestion.

In fig. 8, we show the active monitoring of a link that conversely manifested different levels of congestion. As in the first case this is the monitoring for 24-hour period. For the first part of the day (at night) the throughput maintained the maximum value of 12 Mbps.

Just by looking at the active monitoring we can notice clearly three parts on this plot, one where the throughput is stable, another with a light-medium congestion, and the last part that could be noticed how a heavy congestion. Therefore, we can expect to have a null correlation during the first period, than we should have a high correlation between throughput and fast retransmission and lower one with time outs and an opposite behavior on the late hours. From some point of view this could seem the usual case of an access link. Since we can see that during the night hours, after 00:00 the throughput it is stable at the nominal value, and during the first morning hours we can notice that a congestion manifests and the same continues with more degrading effects, where we notice the heavy congestion.




Fig. 9. Correlation of the passive data for the congestion case, green line correlation throughput with fast retransmit, red line correlation throughput with retransmission cause of time out.

By inspecting the statistics provided by Tstat at the server side, we could confirm this intuition as illustrated in fig. 9, where we show the analysis of the passive data, compared with the throughput of fig. 8, that is reported in log scale in fig. 9. The figure is divided into three window, corresponding to, no, medium and heavy congestion. The blue square are data that we have discarded, since there was not enough information from the statistical point of view. The dots with blue color are from the throughput, in red color are the number of retransmission cause of time out and in green the number of retransmission cause fast retransmit. As we expected, and we can also follow from the plot, the green and red lines (reported in absolute value), during the night both correlations are zero, conversely during the morning hours we start noticing congestion. As long as we have a light congestion only fast packet retransmission is present, and from the correlation point of view, it has a higher value. The more the congestion increases the more the number of segments caused by fast retransmit and time out increase, and when heavy congestion occurs the retransmitted segments due to time out, will be more than those from fast retransmit.

The correlation values show that, in case of light congestion we have a prevalence of the correlation between throughput and fast retransmit and, conversely, in case of heavy congestion we have a higher correlation between throughput and packet time out.

Our intuition concerning the kind of congestion was manually checked, since we could find out that such probe accesses suffered a sort of bottlenecked due to Virtual Leased Line, that was out of the control of the operator.

Fig. 9 suggests us to implement specific algorithms that allows us to generate alarms when some anomalies show up. Such alarms could be important input for network management, and in particular such anomaly detections could be used to make suitable changes in the network configuration, so that part of the traffic could be routed in some less congested links. For an instance other paths (physical or logical) could be switched on following a Software Defined Network (SDN) [12] approach as described in [13].

In particular, we could define suitable thresholds to decide the network operations to be done to maintain a specific level of QoS; in fact, we have seen that the congestion level can be detected by the B decreasing and the increasing of r for R_{FT} and then R_{TO} and in particular the heavy condition is verified when R_{TO} gets higher than R_{FT} .

V. CONCLUSIONS

In this work we have demonstrated the reliability of mPlane approach for SLA verification and certification also on wide geographical area network where server and client can be

located in different Countries. By active and passive measurements in ADSL field tests we have observed that when some unexpected behavior is observed, information is correlated with metrics extracted by Tstat, using domain knowledge to find the root cause of the anomaly. Results show the benefit of complementing active measurements with fine grained details obtained from passive measurements. For instance, we are able to pinpoint congestion at the access link for some lines.

Therefore the analysis about the correlation between active and passive tests can provide important clues to understand the reasons of network degradation and therefore some network changes could be tempted to improve the performance, how in the case of light and heavy congestions.

ACKNOWLEDGMENT

Work carried out in the framework of the FP7 mPlane project.

REFERENCES

- [1] <https://www.misurainternet.it/>.
- [2] L. Rea, E. Mammi, "Italian QoS Monitoring network: impact on SLA control", Networks 2012, Rome, October 16-18 2012.
- [3] <https://www.ict-mplane.eu/public/verification-and-certification-service-level-agreements>
- [4] A. Rufini, M. Mellia, E. Tego, F. Matera, M. Mellia, "Multilevel bandwidth measurements and capacity exploitation in gigabit passive optical networks, IET communications, IET Communications Volume 8, issue 18, December 2014, p. 3357 – 3365.
- [5] <https://www.ict-mplane.eu/sites/default/files/public/public-page/public-deliverables/1341mplane-d56.pdf>, Chapt. 1.7.
- [6] A. Rufini, E. Tego, F. Matera, "Multilevel QoS vs QoE Measurements and Verification of Service Level Agreements" European Conference on Networks and Communications, Bologna 23-29 June 2014.
- [7] S. Traverso, M. Mellia, E. Howalik, S. Raffaglio, A. Fregosi, E. Tego, F. Matera, "Exploiting Hybrid Measurements for Network Troubleshooting" Invited at Networks 2014, Sept, 2014 Madeira (PT)
- [8] <https://www.ict-mplane.eu/sites/default/files/public/public-page/public-deliverables/1342mplane-d63.pdf>
- [9] T. V. Lakshman and U. Madhow, The Performance of TCP/IP for Networks with High Throughput-Delay Products and Random Loss, IEEE/ACM TRANSACTIONS ON NETWORKS, VOL. 5, NO. 3, PG. 336-350, JUNE 1997.
- [10] E. Tego, F. Matera, D. Delbuono, "Experimental investigation on TCP throughput behavior in Optical Fiber Access Networks" in press on Fiber and Integrated Optics.
- [11] www.tstat.org/
- [12] M. Channegowda, R. Nejabati, D. Simeonidou, "Software-defined optical networks technology and infrastructure: Enabling software-defined optical network operations", IEEE/OSA Journal of Optical Communications and Networking, 5, Issue: 10, 2013, pp. A274 - A282
- [13] E. Tego, F. Matera, D. Del Buono, V. Attanasio, "Quality of Service Management based on Software Defined Networking Approach in wide GbE Networks" EuMed Telco 2014 Napoli, Novembre 2014.