



mPlane

an Intelligent Measurement Plane for Future Network and Application Management

ICT FP7-318627

**Dissemination, Exploitation And Standardization
First Report**

Author(s):	POLITO	A. Finamore, M. Mellia, S. Traverso
	FUB	F. Matera, E. Tego
	SSB	G. De Rosa, S. Pentassuglia
	TI	A. Capello, F. Invernizzi
	ALBLF	Z. Ben-Houdi
	EURECOM	P. Michiardi, M. Milanesio
	ENST	D. Rossi
	NEC	M. Dusi (ed.), S. Niccolini
	TID	I. Leontiadis, L. Baltrunas
	FTW	P. Casas
	FHA	R. Winter, M. Faath
	ULg	B. Donnet
	ETH	B. Trammell
	A-LBELL	D. Papadimitriou
	FW	A. Fregosi

Document Number: D7.3
Revision: 1.0
Revision Date: 31 Oct 2014
Deliverable Type: Other
Due Date of Delivery: 31 Oct 2014
Actual Date of Delivery: 31 Oct 2014
Nature of the Deliverable: (R)eport
Dissemination Level: Public

Abstract:

This deliverable describes the dissemination activities and standardization achievements of mPlane during the second year. The document provides detailed dissemination activities undertaken by the project, with the publication of a total of over eighty research papers to journals and various conferences and workshops since the start of the project. Furthermore, the document provides mPlane's standardization efforts, the establishment of links to various working groups and a tutorial to the international delegates. Finally, the document provides a detail report on the software open source release and covers the exploitation activities from all the partners.

Keywords: dissemination, publications, standardization, exploitation, open-source software

Disclaimer

The information, documentation and figures available in this deliverable are written by the mPlane Consortium partners under EC co-financing (project FP7-ICT-318627) and does not necessarily reflect the view of the European Commission.

The information in this document is provided “as is”, and no guarantee or warranty is given that the information is fit for any particular purpose. The user uses the information at its sole risk and liability.

Contents

Disclaimer.....	3
Document change record.....	7
1 Introduction.....	8
2 Dissemination and Standardization activities planned.....	9
3 Dissemination activities.....	12
3.1 Project press release and newsletters	12
3.2 Research papers	12
3.3 Journal paper publications.....	14
3.4 Talks, project presentations, seminars.....	15
3.5 mPlane organized workshops and summer schools	19
3.5.1 Undergraduate course, FHA	19
3.5.2 TMA Summer School, London, UK, April 2014	19
4 Initial Standardization activities.....	22
4.1 IETF activities in Year 2	22
4.1.1 IETF working groups	22
4.1.2 IETF meetings	22
4.1.3 RFCs published.....	23
4.1.4 Drafts in progress	23
4.2 IRTF activities in Year 2	24
4.3 ITU liaison	24
4.3.1 ITU-T SG 13	24
4.3.2 ITU-T SG 12	25
5 Initial record of Participation in concertation and future Internet assembly activities.....	26
6 Open source release.....	27
6.1 Tracebox	27
6.2 fingerprinting	28
6.3 DMFSGD.....	28
6.4 10Gbps+ tools	28
6.5 fastPing.....	29

6.6	repoSim	29
6.7	ECN-Spider	29
6.8	IPFIX	30
6.9	QoF	30
6.10	Schedsim.....	30
6.11	GLIMPSE	31
6.12	mSLAcet active probe	31
6.13	Schedule	31
6.14	ADTool.....	31
6.15	Blockmon.....	32
6.16	Tstat.....	32
6.17	Tstat mplane client.....	33
6.18	tStat-Proxy and Supervisor Implementation.....	33
6.19	RilAnalyzer	33
6.20	Mobile Probe.....	33
6.21	mPlane Nodejs Reference Library.....	34
6.22	mPlane HTTPs Transport	34
6.23	mPlane Components.....	34
6.24	DATI mPlane Proxy	34
7	Exploitation activities.....	35
7.1	Usage of Results.....	35
7.1.1	POLITO	35
7.1.2	FUB.....	35
7.1.3	SSB	35
7.1.4	TI	36
7.1.5	ALBLF	36
7.1.6	EURECOM	36
7.1.7	ENST	36
7.1.8	NEC.....	37
7.1.9	TID	37
7.1.10	FTW.....	37
7.1.11	FHA.....	38
7.1.12	ULg	38
7.1.13	ETH	38
7.1.14	A-LBELL.....	38



7.1.15 FW..... 39

Document change record

Version	Date	Editor(s)	Description
0.1	06 Oct 2014	M. Dusi (NEC)	initial draft
0.9	28 Oct 2014	M. Dusi (NEC)	final draft for review
1.0	31 Oct 2014	M. Dusi (NEC)	final

1 Introduction

This document describes the overall dissemination achievements of mPlane during the second year of the project. The details of the dissemination during the first year are described in deliverable D7.2.

mPlane has put great effort into properly disseminating project results. In summary, dissemination of mPlane results has been very successful. A broad dissemination was achieved by using different channels for making project results available to the public: scientific publications, presentations, and demonstrations at research conferences; presentations and seminars at Ph.D. schools; open source software releases of the mPlane software and tools; and the mPlane website.

mPlane has been very successful in disseminating the project's research results through scientific publications: over forty peer-reviewed, scientific publications have been published with mPlane acknowledgment during the second year (more than eighty in total).

To promote the collaboration across projects and the adoption of mPlane tools, mPlane organized a workshop on its architecture and how to code monitoring probes which are mPlane-compatible during the METRICS bootcamp, which took place between August 25 and September 5, 2014 at the Université Catholique de Louvain, Belgium. The attendance at the workshop was of about 30 people. In summary, the workshop constituted a great occasion for mPlane researchers to disseminate the acquired knowledge and to demonstrate the effectiveness of the mPlane approach. Furthermore, mPlane has been a gold sponsor for the prestigious ACM Internet Measurement Conference (IMC) 2013, and for other conferences that focus on topics strictly related to the mPlane vision, such as TMA 2013, TMA 2014, and TRAC 2014.

Similarly, mPlane dissemination extended to events related to (and jointly organized with) some of the aforementioned workshops. Particularly, mPlane was deeply involved in the organization of the 4th TMA PhD School, which saw mPlane partners as presenters and chair of sessions. The school was the first PhD School ever to have been held in cooperation with ACM SIGCOMM (thanks to ACM funding, the school was entirely free for students, and 15 travel grants were also distributed among the 40+ participants) and it was also the first PhD School to be co-located with the main TMA workshop.

In standardization, mPlane has been active in several IETF working groups, and during IETF meetings, and three RFCs have been published during the second year of the project, and other five drafts are in progress.

mPlane released several tools and software libraries: all these software releases are made available to the public at the project website <http://www.ict-mplane.eu/public/software>. Highlights include the mPlane reference implementation, passive probes and tools to perform analysis on wired and wireless networks, mechanisms to infer the quality of experience of end users and to schedule and process computational tasks on a distributed cluster of machines, and active probes to measure user QoS and to verify Service Level Agreement between ISP and the users.

2 Dissemination and Standardization activities planned

Here, we briefly summarize the overall dissemination and standardization plan of the project:

- Project press release, factsheet and presentation;
- Project website;
- Research papers (aiming at disseminating research results);
- Talks and presentations focused on specific project objectives and tasks;
- Standardization contributions;
- Participation in concertation and Future Internet Assembly (FIA) activities;
- Open source release;
- Exploitation activities.

Table 2.1 shows the plans for dissemination and education activities, while Table 2.2 outlines the standardization activities planned within the mPlane project.

Objective	Approach	Outcome	Specific Actions	Key Performance Indicator
Awareness creation: involving the public	On-line dissemination with public information and news related to the topics.	Reaching a worldwide audience and communities of interest	Project website, including software repositories and collection of tools; brochures and posters.	Website up and running from month two until the end of the project; newsletter issued three times a year and at least three press releases.
R&D community dissemination	Publish the project results in suitable events	Reaching the R&D community involved in related areas of research	Participate in European and international conferences presenting project results and exchanging of know-how	At least ten top-tier conference publications, five journal publications; organization of one IEEE special issue on mPlane-related topics.
Scientific dissemination and Education	Lectures and continued dissemination of project results for under- and post-graduates as well for the scientific community.	Reaching the scientific community, researchers and both under- and post-graduate students	Organize seminars and PhD schools, also with hands-on experiences.	One post-graduate school lasting for one week; one series of Dagstuhl seminars; at least one undergraduate project work.
Industrial dissemination	Venues of different target groups to exchange experiences and advances with respect to the state of the art in the field related to mPlane	Reaching Industry, Regulatory Agencies, Policy makers and interest groups	Organize workshops and ah-hoc meetings.	At least one industrial workshop and one workshop with Regulatory Agencies and policy makers; at least two presentations at operator events.

Table 2.1: Plans for dissemination and Education activities within the mPlane project.

Standards Body	Specific Group Targeted	Details of Contribution
IETF	IPFIX (IP Flow Information Export) - OPS Area	Extensions to the IPFIX protocol to take into account application-layer parameters, and passive and active performance measurements.
	IPPM (IP Performance Metrics) - TSV Area	Extensions to the IPPM framework metrics, possibility to extend IPPM to passive measurements.
	NETCONF (Network Configuration) - OPS Area	Extensions of the NETCONF protocol to take into account measurement probe coordination.
IRTF	NMRG (Network Management Research Group)	Contribute mPlane control interfaces to configure probes. Possibility to generate enough interest to spin-off a working/research group (through a BOF)
	NCRG (Network Complexity Research Group)	Proposals based on mPlane technology that help operators to regain control over their complex networks.
ETSI	Technical Committee Speech Processing, Transmission and Quality Aspects (TC STQ)	Extensions of "User related QoS parameters definitions and measurements" (ETSI EG 202 507-4) to capacity greater than e.g., 20Mbps.
ITU-T	Study Group 12 on Performance and Quality of Service (QoS)	Contribute mPlane metrics definitions and mapping of metrics to user-perceived quality.

Table 2.2: Plans for standardization activities within the mPlane project.

3 Dissemination activities

3.1 Project press release and newsletters

To report on the progress since its start, the project is about to issue an additional press release: the goal is to develop contacts with other associated research groups and projects towards wider cooperative work in the planned research activities. NEC took care of the press release, available on the project website at <http://www.ict-mplane.eu/public/news>.

Newsletters are issued to all subscribers whenever relevant updates are available and are collected and publicly accessible at <https://www.ict-mplane.eu/public/news>.

Finally, the project presentation with details of technical activities and project structure is available on the project website: <http://www.ict-mplane.eu>.

3.2 Research papers

The project partners have been very active in publishing a large number of scientific papers in international conferences and journals to inform the research community about the results achieved in the project.

The list below summarizes the research papers published during the second year in different conferences and workshops. Notably, some of the papers also received distinctions and awards for the presented work (highlights in bold font). Joint collaboration among partners of the project has generated a total of eight published papers.

The full list, together with the papers, is also made publicly available through the project website at <https://www.ict-mplane.eu/publications>.

1. A. Bar, A. Finamore, P. Casas, L. Golab, and M. Mellia, *Large-Scale Network Traffic Monitoring with DBStream, a System for Rolling Big Data Analysis*, IEEE BigData, October 2014. [5]
2. Bocchi, E., M. Mellia, and S. Sarni, *Cloud Storage Service Benchmarking: Methodologies and Experimentations*, 3rd IEEE International Conference on Cloud Networking (IEEE CloudNet 2014), Luxembourg, IEEE, October 2014. [8]
3. P. Casas, A. Sackl, S. Egger, R. Schatz, *Characterizing Microsoft Lync Online in Mobile Networks: a Quality of Experience Perspective*, 3rd IEEE International Conference on Cloud Networking (IEEE CloudNet 2014), Luxembourg, IEEE, October 2014. [11]
4. Traverso, S., E. Tego, E. Kowallik, S. Raffaglio, A. Fregosi, M. Mellia, and F. Matera, *Exploiting Hybrid Measurements for Network Troubleshooting*, IEEE Networks, Funchal, PT, IEEE, September 2014. [42]
5. Fiadino, P., A. D'Alconzo, A. Bär, A. Finamore, P. Casas, *On the Detection of Network Traffic Anomalies in Content Delivery Network Services*, ITC26, Karlskrona, Sweden, September, 2014. [23]
6. Imbrenda, Claudio, Muscariello, Luca and Rossi, Dario, *Analyzing Cacheable Traffic in ISP Access Networks for Micro CDN applications via Content-Centric Networking*. In 1st ACM SIG-

- COMM Conference on Information-Centric Networking (ICN-2014), Paris, France, September 2014. [29]
7. P. Fiadino, M. Schiavone, P. Casas, *Vivisecting WhatsApp through Large-Scale Measurements in Mobile Networks*, SIGCOMM, August 2014. [25]
 8. Dell'Amico, Matteo; Carra, Damiano; Pastorelli, Mario; Michiardi, Pietro, *Revisiting size-based scheduling with estimated job sizes*, in Proc. of IEEE MASCOTS, September 2014. [18]
 9. A. Bär, P. Casas, L. Golab, A. Finamore, *DBStream: an Online Aggregation, Filtering and Processing System for Network Traffic Monitoring*, 5th International Workshop on TRaffic Analysis and Characterization (TRAC), Nicosia, Cyprus, August 2014. [4]
 10. A. D'Alconzo, P. Casas, P. Fiadino, A. Bär, A. Finamore, *Who to Blame when YouTube is not Working? Detecting Anomalies in CDN-Provisioned Services*, 5th International Workshop on TRaffic Analysis and Characterization (TRAC), Nicosia, Cyprus, August 2014. [1]
 11. P. Fiadino, A. D'Alconzo, P. Casas, *Characterizing Web Services Provisioning via CDNs: The Case of Facebook*, 5th International Workshop on TRaffic Analysis and Characterization (TRAC), Nicosia, Cyprus, August 2014. [24]
 12. Andrea Araldo and Dario Rossi, *A per-Application Account of Bufferbloat: Causes and Impact on Users*, 5th International Workshop on TRaffic Analysis and Characterization (TRAC), Nicosia, Cyprus, August 2014, **Best paper award**. [3]
 13. Dusi, M., R. Bifulco, F. Gringoli, and F. Schneider, *Reactive Logic in Software-Defined Networking: Measuring Flow-Table Requirements*, 5th International Workshop on TRaffic Analysis and Characterization (TRAC), Nicosia, Cyprus, August 2014. [22]
 14. L. Cittadini, S. Vissichio, B. Donnet. *On the Quality of BGP Route Collectors for iBGP Policy Inference*. In Proc. IFIP Networking. June 2014. [14]
 15. D.Papadimitriou, *Stochastic Optimal Control in Cooperative Multi-Agent Systems*, Proceedings of ICCSA 2014, Normandie University, Le Havre, France - June 2014.
 16. Rufini A., Tego E., Matera F. *Multilevel QoS vs QoE Measurements and Verification of Service Level Agreements* European Conference on Networks and Communications, Bologna, June 2014. [35]
 17. Casas, P., P. Fiadino, A. Bär, A. D'Alconzo, A. Finamore, and M. Mellia, *YouTube All Around: Characterizing YouTube from Mobile and Fixed-line Network Vantage Points*, EuCNC, Bologna, IT, June 2014. [10]
 18. A. Rufini A., Tego E., Matera F., Mellia M. *Bandwidth Measurements and Capacity Exploitation in Gigabit Passive Optical Networks*, Fotonica 2014, Napoli, May 2014. [36]
 19. Tego E. *Active measurements and limitations of TCP protocol during SLA test* Poster at the Sixth Workshop on Traffic and Monitoring Analysis (TMA), April 2014. [37]
 20. Georges Nassopoulos, Dario Rossi, Francesco Gringoli, Lorenzo Nava, Maurizio Dusi and Pedro Maria Santiago del Rio, *Flow management at multi-Gbps: tradeoffs and lessons learned*. In Traffic Measurement and Analysis (TMA), pages 1-14 , April 2014. [33]

21. Dario Rossi, Guilhem Pujol, Xiao Wang and Fabien Mathieu , *Peeking Through the BitTorrent Seedbox Hosting Ecosystem*. In Traffic Measurement and Analysis (TMA), April 2014. [34]
22. S. Colabrese, D. Rossi and M. Mellia, *Aggregation of statistical data from passive probes: Techniques and best practices*. In Traffic Measurement and Analysis (TMA), pages 38-50 , London, UK, April 2014. [15]
23. Z. Ben Houidi, G. Scavo, S. Ghamri-Doudane, A. Finamore, S. Traverso and M. Mellia, *Gold mining in a River of Internet Content Traffic*, in Proc. 6th International Workshop on Traffic Monitoring and Analysis, TMA April 2014. [6]
24. Trammell, B., D. Gugelmann and N. Brownlee, *Inline Data Integrity Signals for Passive Measurement*, Proceedings of the Sixth International Workshop on Traffic Monitoring and Analysis (TMA 2014), April 2014. [41]
25. S. Colabrese, D. Rossi and M. Mellia, *Scalable accurate consolidation of passively measured statistical data*. In Passive and Active Measurement (PAM), Extended Abstract, pages 262-264, Los Angeles, USA, March 2014. [16]
26. P. Casoria, D. Rossi, J. Auge, Marc-Olivier Buob, T. Friedman and A. Pescape , *Distributed active measurement of Internet queuing delays*. In Passive and Active Measurement (PAM), Extended Abstract, Los Angeles, USA, March 2014. [12]
27. R. Mazloun, M.-O. Buob, J. Auge, B. Baynat, T. Friedman and D. Rossi , *Violation of Interdomain Routing Assumptions*. In Passive and Active Measurement (PAM), Los Angeles, USA, March 2014. [32]
28. P. Casas, P. Fiadino, A. Bär, *Understanding HTTP Traffic and CDN Behavior from the Eyes of a Mobile ISP*, Passive and Active Measurement Conference (PAM), Los Angeles, USA, March 2014. [9]
29. C. Chirichella, D. Rossi, C. Testa, T. Friedman and A. Pescape, *Passive bufferbloat measurement exploiting transport layer information*. In IEEE GLOBECOM, December 2013. [13]
30. A. Araldo, D. Rossi, *Dissecting Bufferbloat: Measurement and Per-Application Breakdown of Queueing Delay*. In ACM CoNEXT'13, Student Workshop, pages 25-26, December 2013. [2]
31. A. Finamore, M. Mellia, Z. Gilani, K. Papagiannaki, V. Erramilli, Y. Grunenberger, "Is There a Case for Mobile Phone Content Pre-staging?", Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies, New York, NY, USA, ACM, 12/2013 (**Best Short Paper Award**) [26]

3.3 Journal paper publications

The list below summarizes the research papers published during the second year in different journals:

1. Ignacio Bermudez, Stefano Traverso, Marco Mellia, Maurizio Munafò, *A Distributed Architecture for the Monitoring of Clouds and CDNs: Applications to Amazon AWS*. In IEEE Transactions on Network and Service Management, 2014, *to appear*. [7]

2. W. Du, Y. Liao, N. Tao, P. Geurts, X. Fu, G. Leduc. *Rating Network Paths for Locality-Aware Overlay Construction and Routing*. In IEEE/ACM Transactions on Networking, 2014 (to appear). [21]
3. G.Ellinas, D.Papadimitriou, J.Rak, D.Staessens, J.Sterbenz, and K.Walkowiak, *Practical issues for the implementation of survivability and recovery techniques in optical networks*, Journal Optical Switching and Networking, vol.14, pp.179-193, 2014. [20]
4. Brian Trammell, Pedro Casas, Dario Rossi, Arian Bar, Zied Ben-Houidi, Ilias Leontiadis, Tivadar Szemethy, and Marco Mellia, *mPlane: an Intelligent Measurement Plane for the Internet*. IEEE Communications Magazine, Special Issue on Monitoring and Troubleshooting Multi-domain Networks using Measurement Federations, 2014. [40]
5. YiXi Gong, Dario Rossi, Claudio Testa, Silvio Valenti and Dave Taht, *Fighting the bufferbloat: on the coexistence of AQM and low priority congestion control* (extended version) . Computer Networks, 60:115 - 128, 2014. [27]
6. Claudio Testa and Dario Rossi, *Delay-based congestion control: Flow vs. BitTorrent swarm perspectives*. Computer Networks, 60:115 - 128, 2014. [38]
7. Trammell, B., *Evolving Transport in the Internet*, IEEE Internet Computing, 2014. [39]
8. Grimaudo, L., M. Mellia, E. Baralis, and R. Keralapura, *SeLeCT: Self-Learning Classifier for Internet Traffic*, IEEE Transactions on Network and Service Management, 2014. [28]
9. Rufini, A., M. Mellia, E. Tego, and F. Matera, *Multilevel Bandwidth Measurements and Capacity Exploitation in Gigabit Passive Optical Networks*, IET Communications, 2014. [35]

3.4 Talks, project presentations, seminars

This section summarizes the talks, presentations and seminars that were held during the second year of the project by all the partners, together with a summary of each. All presentations slides are also available at <http://www.ict-mplane.eu/public/talks>. Beyond the reported talks, each partner made sure to introduce and mention the scope of the project at every occasion they introduced their institution (commercially or in academic circles) and their activities.

Furthermore, mPlane has been a gold sponsor for the prestigious ACM Internet Measurement Conference (IMC) 2013, and for other conferences that focus on topics strictly related to the mPlane vision, such as TMA 2013, TMA 2014, and TRAC 2014.

1. Marco Mellia, *mPlane – an Intelligent Measurement Plane for Future Network and Application Management*, presentation at the LEONE meeting, Torino, October 2014.
Summary: This invited talk represented an opportunity for introducing the mPlane project to the LEONE community. The talk presented the idea behind the project and detailed the system architecture as result of the first year of the project, and then proposed some possible topics for current and future collaborations.
2. Stefano Traverso, *Analysis of Traffic Generated by Clouds and CDNs*, presentation at the Alcatel-Lucent Bell Labs, Dublin, September 2014.

Summary: This invited talk focused on the results we obtained from the analysis of traffic generated by Amazon AWS and presented in the paper “A Distributed Architecture for the Monitoring of Clouds and CDNs: Applications to Amazon AWS”.

3. Marco Mellia, *Management for policy verification*, presentation at the Stakeholders Consultation Workshop - Network Technologies Work Programme 2016-2017 (link is external), Brussels, September 2014.

Summary: This invited talk represented an opportunity for introducing the mPlane project to the FIRE and GENI communities. The talks presented the idea behind the project and detailed the system architecture as result of the first year of the project. In particular, the talk addressed a possible usage of the mPlane architecture for instrumenting experimental testbeds.

4. K. Edeline. *mPlane Architecture and How to Code an mPlane-Complain Application/Proxy*, METRICS Bootcamp (UCL – Belgium), August 2014.

Summary: This talk starts with an overview of the mPlane project, its motivations, its main WPs, its benefits for Internet actors (ISPs, Content providers, Customers, ...) and a few use cases. It continues with a description of the mPlane architecture, its principles and its main components. Then, there is a more detailed description of the mPlane information model, focusing on the classes needed to code an mPlane proxy (Schema, Capability, Specification, Result) and of how to interface an mPlane proxy to mPlane. The talk finishes with examples of existing mPlane proxies and how they are coded.

5. M. Dusi. *Blockmon Architecture and How to Code network monitoring applications*, METRICS Bootcamp (UCL – Belgium), August 2014.

Summary: This talk starts with an overview of the mPlane project, its motivations. It continues with a description of the streaming-platform Blockmon: its main components and how to code high-performance probes and network monitoring applications. The talk finishes with examples and hands-on laboratory session.

6. P. Casas. *Building an Intelligent Measurement Plane for the Internet*, Invited Talk at 5th International Workshop on TRAFFIC Analysis and Characterization (TRAC), Nicosia, Cyprus, August 2014.

Summary: This talk introduced the mPlane project to the TRAC community, providing an overview of the mPlane system, and briefly describing some application use-cases showing the added benefits of the mPlane approach.

7. Z. Ben Houidi, *Crowd Sourced Media Curation Based on HTTP Logs*, LINCIS annual seminar, June 2014.

Summary: The talk described the novel service that results from the use case that ALBLF is developing, together with Polito, in mPlane. LINCIS is the Laboratory of Information, Networking and Communication Sciences in Paris, it gathers researchers from different organizations in the area of Paris.

8. D.Papadimitriou, *Stochastic Optimal Control in Cooperative Multi-Agent Systems*, Proceedings of ICCSA 2014, Normandie University, Le Havre, France - June 23-26, 2014

Summary: In this talk, we consider multi-agent systems where each agent performs a joint task of monitoring time-varying and uncertain events modeled by a stochastic process. Monitoring agents operate in a cooperative mode with partial knowledge/visibility of the global

state of the monitoring system. In this context, Stochastic Optimal Control (SOC) enables to formulate monitoring agents best response as an infinite horizon discounted cost minimization problem and derive the optimal control law that agents have to apply to optimize their value function. The performance of monitoring agents is valued by a global cost function which is an integral of running costs plus an intervention cost, modeling an impulse control.

9. F. Matera, *Role of the QoS measurements in future Next Generation Networks*, University of Cassino, May 2014.

Summary: A review of the main evolution of the future wireline and wireless next generation networks is presented looking at ultra broadband aspects. The role of the QoS measurements in these contexts and the difference among the line capacities offered by the ISP and the QoS perceived by the users. Results from the mPlane project and description of Application SLA (ASLA).

10. Marco Mellia, *mPlane: project and architecture for testbeds integration*, presentation at the FIRE-GENI workshop, Boston, May 2014.

Summary: These three invited talks represented three different opportunities for introducing the mPlane project to the community. The talks presented the idea behind the project and detailed the system architecture as result of the first year of the project.

11. Alessandro Finamore, *Passive inference: Troubleshooting the Cloud with Tstat*, seminar at the 4th TMA phd school, London, April 2014.

Summary: This lecture was composed of two part. The former introduced the mPlane project and its challenges, with particular attention towards passive network monitoring. A few examples of today's Internet services have been considered to investigate the issues related to their passive monitoring. In the latter part of the lesson focused on BigData analytic. In particular, Yahoo! Pig, a modern scripting language for Hadoop, has been introduced to the students, and it was requested to use it to solve some simple exercises using preconfigure virtual machines.

12. K. Edeline, B. Donnet. *Revealing Middleboxes with Tracebox*, 4th PhD School on Traffic Monitoring and Analysis (London – England), April 2014.

Summary: In this talk, we explain how middleboxes are breaking the TCP/IP End-to-End principles, show the latest numbers related to their actual deployment and show several existing path impairment created by middleboxes rewrite or drop policies. Then, we explain in details the tracebox mechanism and how to detect and localize middleboxes using Tracebox. We show early results of a large-scale Tracebox measurements campaign. Participants are asked to code a simplified Tracebox in Python 2 using Scapy and to run it on multiple netkit-emulated network topologies to detect and localize middlebox modifications.

13. G. Leduc, Y. Liao, W. Du, P. Geurts. *Machine Learning-Based Algorithms to Infer End-to-End Network Performance Matrices*, INFORMS Telecommunications (Lisboa – Portugal), March 2014.

Summary: The knowledge of end-to-end network performance metrics is essential to many Internet applications. As active probing of all pairwise paths is unfeasible in large-scale networks, a natural idea is to measure a few pairs and to predict the other ones. We formulate this prediction problem as matrix completion, which is solvable because strong correlations among network path properties exist and cause the constructed matrix to be low rank. The new formulation circumvents the well-known drawbacks of existing approaches based on

Euclidean embedding. In particular, it is applicable to various, possibly asymmetric and non additive, metrics, such as round trip times and bandwidth. It can also be used for rating network paths, i.e., predicting quantized measures of path properties. Compared to fine-grained measurement, coarse-grained ratings are appealing in that they are not only informative but also cheap to obtain. By observing similarities to recommended systems, we show that our inference problem can be solved by a class of matrix factorization techniques.

14. D.Papadimitriou, *Automated analysis and mining of multi-level measurement data*, invited talk at Second NSF Workshop on perfSONAR based Multi-domain Network Performance Measurement and Monitoring (pSW 2014), Arlington (VA), February 20-21, 2014.

Summary: The initial objective of the talk was related to methodology for routing data analysis and detection of (hidden) relationships between dynamic properties (in particular, the stability properties) of routing paths and their corresponding forwarding paths. The method consists in associating traceroute-like record of forwarding paths to the corresponding routing paths derived from the local BGP routing information bases. The resulting tuples recorded over long observation periods are analyzed by means of multi-label classification and non-parametric statistical analysis using non-parametric statistical analysis (e.g. Recurrent Event Data Analysis and Kaplan-Meier estimation).

The talk aimed also to discuss the processing of BGP routing datasets following that the observation that the underlying AS topology is actually becoming a mixed graph because not all AS's are associated to network service providers. However, these datasets are still represented homogeneously and uniformly (and often limited to spatial information) even if they include content service providers and application service providers delivering well beyond transit/connectivity service. In turn, existing topology analysis techniques need to be improved in order to cope with these additional dimensions.

15. M. Dusi, *mPlane - Building an Intelligent Measurement Plane for the Internet*, invited talk at Second NSF Workshop on perfSONAR based Multi-domain Network Performance Measurement and Monitoring (pSW 2014), Arlington (VA), February 20-21, 2014.

Summary: This talk introduced the mPlane project and the addressed problems, i.e., the need of a monitoring plane for the Internet. It further details the system architecture and discusses possible enhancement of current network monitoring systems.

16. F. Matera *Correlation between QoS and QoE measurements: the role of the MPLANE project*, Scuola Superiore delle Telecomunicazioni (Ministero dello Sviluppo Economico), January 2014.

Summary: An overview about the QoS measurement methods is reported with particular details for the ones defined in the framework of the MPLANE project with the description of the use cases. Results on GPON networks are reported. A comparison among QoS and QoE (MOS) measurements is illustrated for GPON e 3G environment.

17. Y. Vanaubel, B. Donnet. *Network Fingerprinting: TTL-Based Router Signatures*, Telecom Paris-Tech (Paris – France), March 2013, and Queen Mary University of London (London – United Kingdom, November 2013).

Summary: Fingerprinting networking equipment has many potential applications and benefits in network management and security. More generally, it is useful for the understanding of network structures and their behaviors. In this paper, we describe a simple fingerprinting mechanism based on the initial TTL values used by routers to reply to various probing messages. We show that main classes obtained using this simple mechanism are meaningful

to distinguish routers platforms. Besides, it comes at a very low additional cost compared to standard active topology discovery measurements. As a proof of concept, we apply our method to gain more insight on the behavior of MPLS routers and to, thus, more accurately quantify their visible/invisible deployment.

18. Rolf Winter, *Internet Research and Education* Invited Guest Speaker at the RISE workshop organized by the Ritsumeikan University in Kyoto, Japan, October 2013.

Summary: The talk focused mainly on the mPlane project goals and the work carried out at the FHA. It introduced the GLIMPSE probe and the structure of the back-end and also showed the standardization aspects of the project.

Together with students from the design faculty, the GLIMPSE team produced an informational video (published on Youtube), that is intended to explain the project idea to non-technical audiences. The Glimpse video will be used as a vehicle to motivate non-technical people to install the app.

The GLIMPSE team has set up various social channels for the GLIMPSE probe including Twitter, G+, Facebook etc. They will be used to disseminate project progress and updates to the GLIMPSE users.

3.5 mPlane organized workshops and summer schools

3.5.1 Undergraduate course, FHA

At FHA an undergraduate course has been organized where a large student team helped develop the GLIMPSE probe and back-end. It worked on various aspects of the mPlane project.

3.5.2 TMA Summer School, London, UK, April 2014

ENST chaired the organization of the 4th PhD school on Traffic Monitoring and Analysis (TMA).

This 4th school marks three nice important achievements:

- it was the first PhD School ever to have been held in cooperation with ACM SIGCOMM; thanks to ACM funding, the school was entirely free for students, and 15 travel grants were also distributed among participants
- it was the first PhD School to be co-located with the main Traffic Monitoring and Analysis (TMA) Workshop, now at its 6th edition. Students of the School were entitled following the TMA workshop at no charge.
- all the material, including video lectures, virtual machines, instructions and solutions are available online – extending the reach beyond that of the about 40 participants.

The school was deeply connected to mPlane, with intervention by ULg and POLITO, disseminating achievements obtained during the course of the project, in such a way that students would be able to replicate their methodology

Following is an excerpt of the summary of school activity and organization from <http://networks.eecs.qmul.ac.uk/news/tma-2014/phdschool/>. For more details, please refer to the URL.

3.5.2.1 At a glance

School

- 4th PhD School on Traffic Monitoring and Analysis (TMA), 2014
- in cooperation with ACM SIGCOMM
- <http://networks.eecs.qmul.ac.uk/news/tma-2014>

Organizers

- Dario Rossi, Chair (tma2014-phd-school@listes.telecom-paristech.fr)
- Sabri Zaman, Local chair (sabri.zaman@qmul.ac.uk)
- Hamed Saljooghinejad, Local chair & Webmaster (h.saljooghinejad@qmul.ac.uk)

Speakers

- Jordan Auge and Marc-Olivier Buob (Université Pierre et Marie Curie)
- Benoit Donnet and Korian Edeline (Université de Liege)
- Renata Teixeira and Anna-Kaisa Pitilainen (Inria)
- Alessandro Finamore (Politecnico di Torino)

Dates and venue

- April 15th-16th, London UK
- Colocated with the 6th TMA Workshop, held the April 14th
- Room: People2019s Palace 2 (PP2), School of EECS at Queen Mary University of London, Mile End Road, London E1 4NS, United Kindom

3.5.2.2 Summary

The 4th PhD school on Traffic Monitoring and Analysis (TMA), continues the traditional blend of theory and practice initiated by the three former editions (see the TMA portal <http://www.tma-portal.eu/cost-tma-action/phd-schools/>). Talks were followed by hands on laboratory sessions where PhD students had the chance to put in practice cutting-edge methodologies they just have been exposed to, aiming at reinforcing the learning process and drive down the cost to start new research work in the TMA domain. This year school covered both passive and active measurement techniques, with emphasis on inference of network behavior, possibly at scale, covering the whole value chain of the Internet ecosystem.

In the current Internet, users access Cloud based applications through Web browsers or smart-phone apps. The overall user experience depends on different and independently operated network segments and software tools: so that the overall path from the household to the Cloud starts from an home network and traverse several autonomous systems, before finally hitting the data center or content distribution node. It follows that to assess the quality of user experience, and troubleshoot issues that potentially arise at any step in the path, a multitude of methodologies are needed.

For example, problems may be local to the user home (e.g., bufferbloat), or in the path (e.g., failing or misconfigured equipment), or in the datacenter (e.g., faraway Cloud node). Techniques and

tools learned throughout the school will allow students to detect Cloud applications that are using anycast addresses by leveraging multiple vantage points. They are able to reveal and fingerprint intermediate routers in the paths to the Cloud. They are able to assess application performance from the edge, either by peeking through the user browser, or by passive inspection of the traffic from network links.

The PhD school took place right after the TMA conference and school attendees had the chance to follow the TMA workshop at no cost. Additionally, by means of poster sessions, gave PhD students a chance to present their ongoing work to TMA attendees and took valuable feedback. Finally, participants could share her/his research with peers and specialized professors during practical sessions with the aim of advancing her/his own project and consolidating a network of practitioners and scholars in the TMA field.

4 Initial Standardization activities

This Section describes the standardization achievements of the mPlane project during the second year. The standardization strategy taken is outlined and the concrete activities in the IETF, IRTF and ITU are described.

4.1 IETF activities in Year 2

4.1.1 IETF working groups

IETF LMAP Working Group – ETH contributes to discussions in the LMAP working group and is monitoring the working group for opportunities to contribute mPlane work to LMAP. Michael Faath (FHG) has attended the meeting, with discussion with partners of the EU Leone projects to synchronize on the LMAP protocol implementation.

IETF IPPM Working Group – Brian Trammell from ETH co-chairs the IP Performance Metrics (IPPM) working group; as an IPPM chair, he is overseeing metric development for large-scale network measurements. Where possible, mPlane will use IPPM measurement protocols and metrics for active measurement.

IETF IE-DOCTORS group of experts – Brian Trammell from ETH is a member of the IE-DOCTORS group of experts responsible for the maintenance of the IPFIX Information Element Registry, on which work on the mPlane Element Registry is based.

IETF tcpm Working Group – The TCP Maintenance and Minor Extensions (tcpm) working group handles small changes to TCP. ETH is working on a more detailed Explicit Congestion Notification (ECN) feedback that is an active work item in tcpm. ECN is a TCP/IP extension to signal congestion with the introduction of loss. Further the redefinition of ECN semantics to promote its wide-scale deployment is under discussion in tcpm with contribution of ETH. A wide-scale deployment of ECN would provide more fine-grained information on congestion occurred previously on the path that can also be used for passive measurements.

IETF ConEx Working Group – The Congestion Exposure (ConEx) working group specifies an IP signaling protocol that exposes congestion information known by the end hosts to the network. With the ConEx protocol whole-path congestion information can easily be assessed by each point on the network path. This information can be used for traffic management and passive measurement of the current congestion level (over different time scales). Mirja Kühlewind is a co-author of two working group documents. The ConEx working group is currently finalizing the standardization process.

4.1.2 IETF meetings

Partners participated in the following IETF meetings during the period:

IETF 88 Vancouver, Canada – Brian Trammell attended the IETF 88 meeting in Vancouver, where he co-chaired the IPPM Working Group meeting, and attended side-meetings relevant to the

creation of a framework for passive metrics within IPPM. Michael Faath participated in the IPPM and LMAP working groups.

IETF 89 London, UK – Brian Trammell attended the IETF 89 meeting in London, where he was appointed to a two-year term as a member of the Internet Architecture Board (IAB), co-chaired the IPPM Working Group meeting, and co-chaired the Tunneling Compressed Multiplexed Traffic Flows (TCMTF) Birds-of-a-Feather (BoF) session. He also attended sessions of interest to mPlane, especially with respect to passive measurements of transport performance: TAPS, TCPM, RMCAT, and LMAP, as well as a side meeting to coordinate developments with the LEONE project, where also Michael Faath attended.

IETF 90 Toronto, Canada – Brian Trammell attended the IETF 90 meeting in Toronto, where he chaired the first meetings of the IAB’s IP Stack Evolution Program and co-chaired the IPPM Working Group meeting. He also attended sessions of interest to mPlane, especially with respect to passive measurement of transport performance: TSVWG, TCPM, TAPS, and RMCAT, as well as side meetings to discuss future deployments of QoF. Michael Faath attended the meeting as well, where he focused on LMAP and IPPM WG participation.

4.1.3 RFCs published

Here follows a list of the RFCs being published during the second year of the project:

1. RFC 7125 – Trammell, B. and P. Aitken, “Revision of the tcpControlBits IP Flow Information Export (IPFIX) Information Element”, 02/2014. This RFC expands the ability to do passive TCP observation using IPFIX by including TCP flags information relevant to explicit congestion notification (ECN); the use of ECN, in addition to reducing performance penalties due to congestion, improves passive observability and localization of performance issues as
2. RFC 7119 – Claise, B., A. Kobayashi and B. Trammell, “Operation of the IP Flow Information Export (IPFIX) Protocol on IPFIX Mediators”, 02/2014. Defines extensions to the IPFIX protocol for use on mediators; i.e., any mPlane repository component both accepting and re-exporting IPFIX.
3. RFC 7373 – Trammell, B., “Textual Representation of IP Flow Information Export (IPFIX) Abstract Data Types”, 09/2014. Defines textual representations for IPFIX primitive types to be compatible with the JSON representations used for equivalent primitives in mPlane (see D1.4 section 4.1.1).

4.1.4 Drafts in progress

The following contributions are currently in progress; all drafts are focused on the deployment and usability of ECN, thereby improving passive observability of congestion- and queueing-related TCP performance issues.

1. “A Mechanism for ECN Path Probing and Fallback” (draft-kuehlewind-tcpm-ecn-fallback-01) ECN is hardly used due to the perceived unusability of ECN on many paths through the Internet caused by ECN-ignorant routers and middleboxes. This document specifies an ECN

probing and fall-back mechanism in case ECN has been successfully negotiated between two connection endpoints, but might not be usable on the path. A paper that further assesses the status of ECN deployment on web servers and connection failures link to ECN negotiation attempts is currently under submission.

2. “Problem Statement and Requirements for a More Accurate ECN Feedback” (draft-ietf-tcpm-accecn-reqs-07)
This document is a working group item in the tcpm working group and describes requirements for a more accurate ECN feedback. A more accurate ECN feedback is needed for the ConEx signaling protocol as well as for a more fine-grained congestion response.
3. “More Accurate ECN Feedback in TCP” (draft-kuehlewind-tcpm-accurate-ecn-03)
This document proposes a concrete new ECN feedback mechanism to provide more accurate ECN feedback. ETH plans to implement the proposed mechanism in Linux for further evaluation.
4. “IPv6 Destination Option for ConEx” (draft-ietf-conex-destopt-06)
This document describes the specification of the ConEx signaling protocol as IPv6 ConEx Destination Option (CDO). The ConEx congestion information is provided by the sender and can be read by each node on the network path to estimate the congestion level on the path. Therefore ConEx can be used for passive congestion measurements.
5. “TCP modifications for Congestion Exposure” (draft-kuehlewind-conex-tcp-modifications-01)
This document specifies actions needed by a TCP sender to provide the ConEx congestion information in the CDO.

4.2 IRTF activities in Year 2

A-LBELL actively contributes to the NMRG (Network Management Research Group) of the IRTF.

Dimitri Papadimitriou from A-LBELL has participated to the NMRG/IRTF meeting held in London on March 14 2014 with a presentation on the problem of probe placement.

4.3 ITU liaison

FUB supports ISCOM (Ministero dello Sviluppo Economico), that is the Italian ITU representative, about topics of the SG12 (QoS), SG13 (SDN), SG15 (optical networks). To this end, FUB shared with ISCOM the topics being investigated in mPlane regarding the QoS and QoE measurements (SG12), and the control plane in SDN networks driven by QoS measurements.

4.3.1 ITU-T SG 13

FUB has been working on a ITU-T normative based on the approach described in the paper “Quality of Service Management based on Software Defined Networking Approach in wide GbE Networks” by E. Tego *et al.*: the approach focuses on the network management driven by QoS active and passive tests. The main idea is based on the fact that the mPlane reasoner could be one of the input of the SDN orchestrator.

4.3.2 ITU-T SG 12

Fastweb participated to the kickoff meeting for the establishment of an Italian Commission with the goal of contributing to the ITU Study Group 12 (SG12 - Performance, QoS and QoE). The newly established Commission appreciated the participation of mPlane partners (FW, POLITO, FUB) and their offer to share the mPlane knowledge experience into the commission studies.

During the same meeting, FUB showed the interest in submitting a normative on QoS measurement that is able to distinguish between line capacity (given by the ISP) and available user bandwidth (throughput and goodput). Currently most of QoS tests are still based on ETSI EG 202 057-4, "User related QoS parameter definitions and measurement", October 2005, that is quite obsolete and not suitable for current and future networks.

5 Initial record of Participation in concertation and future Internet assembly activities

The project is aware of a number of concertation activities that the EC regularly organizes to coordinate the research results across projects. The mPlane project participated in the following concertation and future Internet assembly activities:

- “Management for policy verification”, presentation at the Stakeholders Consultation Workshop - Network Technologies Work Programme 2016-2017 (link is external), Brussels, September 2014.
- “mPlane: project and architecture for testbeds integration”, presentation at the FIRE-GENI workshop, Boston, May 2014.

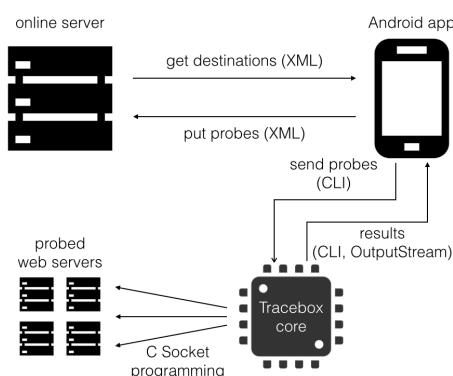


Figure 6.1: TraceboxAndroid architecture overview

6 Open source release

mPlane has been developing quite a number of tools for the analysis and measurement of the Internet infrastructure. The project has released most of those tools as open source to the community and partners are actively involved in the development of them.

6.1 Tracebox

tracebox [19] is a tool for revealing the presence of middleboxes (i.e., “an intermediary box performing functions apart from normal, standard functions of an IP router on the data path between a source host and destination host”) along a path. tracebox uses the same incremental approach as traceroute, it sends packets with different IP, UDP or TCP fields and options with increasing TTL values. By comparing the quoted packet to the original, one can highlight the modifications and the initial TTL value allows us to localize the two or more hops between which the change took place.

We ported tracebox into scamper [31], an all-around parallelized topology and performance analysis tool which implements various simple to complex measurement methods (ping, traceroute, alias resolution, and more). The porting consisted in the previously described middlebox detection mechanism and other techniques. For instance, the proxy detection algorithm performs a first traceroute round over TCP, a second over UDP and compares the number of hops. The statefull box detection algorithm is more complex and relies on multiple TCP traceroute-like probing with decreasing Initial Sequence Number values. It assumes that a statefull box will perform in-window sanity tests and expects specifically wrong packets to be dropped.

scamper has been fully integrated into the mPlane architecture.

We also ported tracebox into Android. Fig. 6.1 illustrates the tracebox Android architecture. We see that it is made of three components: the backoffice (or server – coded in PHP and HTML) that is used to collect data (and make later analysis), the front-office (or the application) corresponding to the Android application (coded in Java) and the system core where the tracebox intelligence has been put (coded in C, under the frontend, and based on busybox, i.e., a smaller versions of common GNU tools, programs, and scripts for embedded systems). The frontend communicates with the server using an XML API that gives the application the destinations to be probed and allows it to

send the results from probes. The application communicates with the Core using an API provided by the Android SDK (using Process and Runtime classes). The core itself implements tracebox (as described above) and sends probes to the destinations using sockets by system calls.

The Android application is available on the Google Play Store since mid-2014 and widely deployed in order to get data from all over the world and the more cellular carriers possible. Additional information can be found at <http://www.androidtracebox.org>.

6.2 fingerprinting

Fingerprinting [43] refers to the act of dividing network equipment into disjoint classes by analyzing messages sent by that equipment, usually in response to some form of active probing. Those classes may correspond, for instance, to router operating system (OS), router brand, or router configuration. We propose a fingerprinting method that is a companion to traceroute-like exploration. Our fingerprinting method infers several values such as the initial TTL values used by routers when building their different kinds of reply packets, ICMP packet sizes, ToS, and DSCP. We call this set of values outer signatures. Router signatures are meaningful for fingerprinting as the initial TTL values (for instance) vary not only between different router platforms but also depending on the protocol and the type of message (error versus standard replies for instance). Indeed, no specific default value has been standardized for the TTL field.

We have ported our fingerprinting techniques into scamper, which has been fully implemented in the mPlane architecture.

6.3 DMFSGD

The DMFSGD (Decentralized Matrix Factorization by Stochastic Gradient Descent) [30] approach is suitable to estimate the properties of all paths, such as round-trip time (RTT) and available bandwidth (ABW), in a large-scale system, without having to measure them all, which would require $O(n^2)$ measurement cost, where n is the number of nodes. To address this issue, a natural idea is network inference whereby only a subset of paths are actually measured while all others are predicted. Although less accurate compared to the measurement of all paths, this “measure a few and predict many” framework is much more scalable due to the significant reduction of measurement overheads.

We have implemented DMFSGD in Python. This implementation is available here: <http://queen.run.montefiore.ulg.ac.be/~liao/DMFSGD.html>. This implementation can be directly deployed on the PlanetLab testbed.

6.4 10Gbps+ tools

We have released software able to generate (and operate) on traffic at 10Gbps+ speed on off-the-shelf hardware. This effort stem from our work on high performance traffic classification [17] completed during Y1 of mPlane, and continued in the [33] work during Y2.

While our main goal in [17] and [33] was to design a high performance traffic classification and flow management engines respectively, however to stress test the system we needed an efficient

packet injection tool, able to operate at 10Gbps and beyond speed, on common hardware.

As our High-speed Pcap (HPCAP) traffic injection tool is useful beyond the focus of traffic classification and management [17, 33], we release them as open source software to the community.

HPCAP is able to send both synthetic and real pcap traces up to 14.2 Mpps and 10 Gbps. Note that CPU and memory affinity is crucial (to avoid or reduce cache misses) in order to get line-rate. Please refer to our papers for more technical details.

6.5 fastPing

Fastping, is a fast ICMP scanner developed during the Y1 of mPlane. Fastping can be run on a off the shelf PC, with a probing rate of about 50k hosts per 5 seconds. Scalability of a single probe is obtained in user-space (as opposite to the zmap software that requires root privileges), with a non-blocking multi-thread design (that allows to significantly exceed nMap Scripting Engine performance, but of course not as much as zmap).

While we released the tool as open source at <http://www.ict-mplane.eu/public/fastping>, the benefit of it comes from leveraging a distributed measurement infrastructure such as PlanetLab. In Y2 we have demonstrated the usefulness of Fastping by large-scale bufferbloat measurement [PAM-14b].

Currently, we are building an Internet-scale anycast scanner on top of Fastping (publication under submission), that we plan to release not as a tool (as it would be a weapon for DDoS) but rather as a service (possibly in form of a safer mPlane repository).

6.6 repoSim

repoSim is an ns2-based simulator aimed at assisting the fine-tuning of mPlane repository performance. The overall goal would be to use simulation as a preliminary, necessary step to investigate a broad spectrum of solutions, to find candidate solutions worth implementing in real operational mPlane repositories.

While we have released the tool as open source, for the time being repoSim is likely useful only internally to project partners <http://www.ict-mplane.eu/public/reposim>.

6.7 ECN-Spider

ETH has released an ECN probing tool in the frame of the project available under the GNU Lesser General Public License (<http://ecn.ethz.ch/ecnspider>). This tool is complementary to Tracebox, and part of a larger effort to scalably build maps of path impairment on the Internet (e.g., paths which exhibit traffic- or header-specific connectivity or performance issues) as background information for queries about connectivity issues stored at a repository.

Based on this tool ETH has performed ECN measurements to document the state of ECN deployment and connection failures link to ECN negotiation attempts on the top million websites according to Alexa. A respective paper is under submission and measurement results are available online (<http://ecn.ethz.ch>).

6.8 IPFIX

ETH is continuing the development of the ipfix module for python 3.3 (<http://pypi.python.org/ipfix>) available under the GNU General Public License from the public GitHub repository (<http://github.com/britram/python-ipfix>). The IPFIX module makes it simple to build IPFIX collecting, exporting, and intermedial processes in Python, and is intended within mPlane to provide a low-volume bridge between mPlane result messages and IPFIX messages.

6.9 QoF

ETH is continuing the development of the QoF TCP-aware flow meter based on YAF available under the GNU General Public License from the public GitHub repository <http://github.com/britram/qof>. QoF is an IPFIX Metering and Exporting process, designed for passive measurement of per-flow performance characteristics. It is primarily intended to support research into passive measurement of performance metrics for TCP flows; however, it can also be used for general flow measurement, especially in environments where the deployment of technologies which inspect packet payload is restricted. QoF is a fork of YAF version 2.3.2, with the following major differences from the YAF codebase:

- Removal of all payload inspection code.
- Replacement of packet acquisition layer with WAND's libtrace.
- Replacement of most command line flags with a YAML-based configuration file, which allows implicit feature selection through direct specification of the information elements to appear in QoF's export templates; this configuration file is designed to be easily mapped to a corresponding mPlane capability message.
- Support for new information elements focused on passive TCP performance measurement.

6.10 Schedsim

This is a simulator for evaluating the impact of errors in estimating the size when performing size-based scheduling in big-data workloads. Details in our technical report, available at <http://arxiv.org/abs/1306.6023>, while the source code is available here: <https://github.com/DistributedSystemsGroup/schedsim>.

Needed software:

- wget (to get the datasets)
- Python 3.2 or 2.7
- Python libraries: numpy, matplotlib (for plots), blist

Pay attention: blist v1.3.4 has a trivial bug (a couple of missing underscores) that makes it not work in Python 3.2. At the time of this writing, the problem is fixed on github, but not on the version that can be installed from pypi. If you want to manually fix the bug, it's sufficient

to replace `self._mapping.sortedkeys` with `self._mapping._sortedkeys` in lines 32 and 53 of `_sorteddict.py`.

6.11 GLIMPSE

GLIMPSE is an end-host based probe and the respective infrastructure (repository, supervisor and reasoner), which enables true end-to-end measurements plus measurement out of home networks. It has been developed cross platform - running (at least) on Linux, Windows, Mac OS and Android. Both GUI and console versions will be provided. The app runs in the same spirit as SETI@home where CPU time is donated. Instead of CPU time, the GLIMPSE probes rely on bandwidth donations. Common measurement methods are implemented (Ping, traceroute, Speedtest etc.) - all without requiring root privileges so that rooting of devices is not necessary. Its primary use case is network troubleshooting.

A GLIMPSE probe release has been published through the project website at <http://www.ict-mplane.eu/public/software>. The current software development is done privately until the first production-stable release. After that, all software code will be made available through github.

6.12 mSLAcet active probe

mSLAcet active probe was upload in <https://www.ict-mplane.eu/public/mslacert-active-probe>, together with instructions on how to run it. The official version is available at <https://github.com/etego/msla>.

The probe is based on an algorithm that is capable to give mSLA certification, by making use of both the UDP or the TCP protocols. To achieve the mSLA certification the algorithm makes use of the tool iperf (<http://iperf.fr>), and PING. mSLAcert makes the measurement and calculation to certify, the good-put at Layer-7, throughput at Layer-4 and line capacity at Layer-2 of the OSI stack.

6.13 Schedule

The tool Schedule is a tool for cache-oblivious scheduling of shared workloads in big-data analysis repositories. Schedule can be used to generate schedules of a given set of analysis jobs, solving all precedence constraints among inter-dependent jobs and, depending on the selected algorithm, optimizing the complete operation for minimum cache usage. Schedule is currently running on top of the DBStream repository.

6.14 ADTool

ADTool is a software tool that implements an Statistical Anomaly Detection Algorithm capable of detecting network and traffic anomalies in near real-time, running on top of DBStream. The software requires suitable DBStream jobs to compute traffic feature distributions with the required time-granularity. It is designed to run in near real-time, processing the distributions of the moni-

tored features as soon as they are available in DBStream.

6.15 Blockmon

NEC is continuously developing and maintaining Blockmon, an Open Source stream-processing platform, designed to enable measurement of ever larger networks which is available on the project website and at the following two GitHub repositories: the Blockmon node (<https://github.com/blockmon/blockmon>) and the Blockmon controller (<https://github.com/mdusi/blockmon-controller>).

The platform now is able to run distributed applications and relaunch the computation in case of node failure. It is actively under development the implementation of primitives for stream-line flow processing, able to keep track of a large number of flows with a limited memory footprint. Evaluation of the technique is ongoing.

6.16 Tstat

Tstat is an open source traffic passive analyzer developed at Telecommunication Network Group (TNG) (link is external) of Politecnico di Torino (link is external). It started as a branch of TCPtrace (link is external) with a focus on TCP statistics only, but over the years it evolved in a full fledge monitoring solution offering an extensive set of measurements for IP, TPC and UDP, as well as traffic classification capabilities through a combination of Finite State Machine Deep Packet Inspection (FSM DPI) and Behavioural engines.

Thanks to the support of the mPlane project we extended Tstat functionality with the following features

- HTTP module: it allows to save text log files reporting information about HTTP queries/responses
- IP address anonymization: it allows to mask local IP address monitored using hashing functions
- Blockmon integration: we collaborated with NEC to integrate Tstat analysis modules in Blockmon
- log_sync: a client/server application which allows to continuously export from Tstat logs from probes to repositories
- improved log configurability: rather than collecting a monolithic set of stats, Tstat now offers more fine-grained control on which set of features a saved in the logs. Per-flow stats are now grouped in macro classes which can be added or removed at runtime
- Android integration: thanks to the effort of TID, Tstat works also on rooted Android devices
- OpenWRT integration: thanks to the effort of NETVISOR, starting from release 37196, the OpenWRT Linux distribution contains Tstat as a "Network Utilities" package

6.17 Tstat mplane client

Beside the efforts related to the evolution of Tstat, we also developed a “proxy”, i.e., a mPlane client capable of interfacing Tstat with the mPlane system. This software have been release and documented as part of D2.2.

6.18 tStat-Proxy and Supervisor Implementation

Python implementation derived from the original Reference Implementation, with several additions and modifications. The message exchange approach has been modified from capability pull, specification push, to capability push, specification pull. A Supervisor has been implemented, that works as an HTTP server: now, all the components interact only through it. HTTPS has been adopted as standard, hence no more HTTP is supported. (<https://github.com/stepenta/RI>)

6.19 RilAnalyzer

RILAnalyzer: a tool to perform network analysis from within a mobile device.

Modern smartphone platforms add new challenges for the cellular networks they are running on. Equally though, complexity of cellular networks is hard to deal with for application and system developers, worsening each others performance and efficiency. Such difficulties are largely caused by the lack of cross-layer understanding of interactions between different entities - applications, devices, the network and its management plane.

To address the issue, we are releasing RILAnalyzer publicly. It is a tool that provides mechanisms to perform network analysis from within a mobile device. RILAnalyzer is capable of recording low-level radio information and accurate cellular network control-plane data, as well as user-plane data. Such data can be used to identify previously overlooked network and connectivity management issues and infer how the different configurations interact with application logic, causing network and energy overheads.

6.20 Mobile Probe

The application periodically launches a YouTube video from a list of short and long videos, while it logs hardware, system and network measurements in the background. When a number of measurements becomes available, they are aggregated and sent to a remote server. For the network measurements the application uses a precompiled tstat binary for Android which is packed inside the apk file.

Collected Metrics:

Hardware: CPU usage, free memory, RSSI (when on WiFi), cell tower information, location information from GPS and WiFi, connectivity state
System: Playback state, re-buffering events, re-buffering duration, load time, HTTP requests, video decoder state
Network: statistics per tcp flow as they are provided by tstat
Requirements: The application requires root access and the pcap library to be present in /system/lib in order to allow tstat to listen on the network interface. It is also required

to have the official YouTube application installed on the device.

6.21 mPlane Nodejs Reference Library

The mPlane Reference Library has been implemented in Nodejs. The architecture and software structure is freely inspired by mPlane reference implementation written in python by Brian Trammell. The library is fully compatible with the Python implementation and is available through the NPM package manager.

6.22 mPlane HTTPs Transport

This library implements a Nodejs module for transport of mPlane informational elements over HTTPS. All transaction can be secured using trusted certificates.

6.23 mPlane Components

This package contains working examples of three mPlane components: a Supervisor, a probe (ping and traceroute) and a simple client for the supervisor. The implementation leverages mPlane Nodejs library and mPlane HTTPS transport library. A complete set of working SSL certificates is provided (with root-CA and signing-CA) in order to have a complete, full working environment.

6.24 DATI mPlane Proxy

The mPlane proxy interface for DATI is written in Nodejs leveraging on the mPlane and HTTPs Nodejs libraries.

7 Exploitation activities

mPlane explores new areas for monitoring applications by research and prototyping. However, the outcome of mPlane is also intended to be used by each partner to advance in their particular area. This section is describing the usage of results stemming from the mPlane project.

To the purpose of exploiting the project as a whole, together with the University of Auckland, one of the collaborating institutions, we have published a paper on *Inline data integrity signals for passive measurement* in LNCS 2014. Moreover, the University of Auckland has completely re-worked the outer edge of their campus network, and upgraded their research measurement infrastructure so that it is possible now to collect packet traces from a 10Gb/s interface. The institution would be open to sharing trace data with other mPlane researchers.

7.1 Usage of Results

Based on the initial achievements accomplished during the first two years of the project, each partner has set the basis for a good exploitation of mPlane results. Here follows the relevant factors and the initial usage of results from each partner, together with plans on how to continue making use of such results.

7.1.1 POLITO

POLITO is already currently exploiting some of the functionalities of mPlane to monitor the Politecnico di Torino campus network. Similar activities are carried on also in collaboration with FW to monitor a real ISP network. Such activities will be continued, and possibly extended to other partners.

7.1.2 FUB

FUB, with its support to the Italian Communication Regulation Body (AGCOM), will try to use the mSLAcert probe, just implemented in the framework of mPlane, to introduce a user agent to verify the SLA between ISP and client in terms of bandwidth.

7.1.3 SSB

SSB will take advantage of the knowledge and the code developed so far within the project, creating an infrastructure that will be used to monitor the performance of distributed and data streaming applications of its commercial partners and promptly intervene in case of problems. In particular, SSB has started to enrich FID framework with new QoS monitoring features and troubleshooting improvements. FID is a framework developed by SSB for banking market, to manage and deliver real time big financial data originating from international trading markets. SSB is discussing with a primary italian bank to introduce the new mPlane-based monitoring and troubleshooting features in their internal financial information delivery infrastructure. SSB plans to start a trial in the next months.

7.1.4 TI

TI is interested in developing and deploying a measurement infrastructure able to integrate different probes, developed internally or bought from third parties, and to build an intelligent system to use data collected by the probes to enhance network operations. The results of the project will serve as a basis over which building this measurement infrastructure. The current status and the achievements of the mPlane project have been shared and discussed within Telecom Italia Technology Plan, which is the document that drives TI's innovation activities for the following three years.

7.1.5 ALBLF

ALBLF intends to push further the use case on which it is working, by continuing building intellectual property around it, and by looking for means of "transfer". This use case demonstrates how the mPlane architecture and tools can be instrumented to offer novel added value services, like data analytic reports and media curation platforms that recommend interesting content to users, based on the analysis of network content traffic. ALBLF continued building intellectual property around this topic.

Thanks to mPlane, ALBLF worked on building a prototype of a service, called NetCurator, that leverages the network content traffic to detect trending topics and interesting content, and promote them to the users. In collaboration with Polito, this prototype is running today in the campus network which will ultimately help us (1) enhancing the service and (2) understanding what the users think about it. The first users of NetCurator left encouraging feedback. As a next step, we intend to deploy NetCurator in an enterprise network, to understand whether it is interesting or not to have it in such a different environment/community.

7.1.6 EURECOM

EURECOM plans to use the output and key concepts developed in mPlane to enrich the current Master-level and PhD-level curriculum for students attending EURECOM. In particular, all work related to scheduling will be covered in the *Distributed systems and Cloud Computing* course. Several semester projects have been and will be dedicated to work on the mPlane testbed at Eurecom, including the design of better data transport mechanisms and additional data analysis algorithms not covered in the DoW of mPlane nor in the current use-cases.

7.1.7 ENST

ENST plans to exploit the scientific results obtained in mPlane on both educational and research directions. As far as education is concerned, (i) ENST already chaired a PhD school on the subject, (ii) some of the activities carried on during mPlane are already exposed to students at Ecole Polytechnique (projects during INF570 course) and at Telecom ParisTech (seminar during RES224 course), and (iii) a more complete course on Internet measurement will be possibly launched at Ecole Polytechnique next year, which will leverage on the body of knowledge gained during the mPlane project

As far as the research is concerned, ENST plans to increase its international notoriety in the field of high speed packet processing (beyond 10Gbps), and continuous large-scale Internet topology discovery, activities that were initiated during mPlane. Notoriety is hopefully acquired through both publications in reference venues in the field (ACM SIGCOMM IMC, Passive and Active Measurement), distinctions of such publication (1 best paper award during Y2 and 1 runner up for best paper during Y1), organization of related events (1st ACM SIGCOMM PhD School) as well as through the release of open source software included in the mPlane suite (e.g., the fastPing or 10Gbps+ tools described above).

Notoriety can also possibly be reached by implementing and running constantly successful distributed services. Notably, the effort on fastPing during Y2 led to the proposal of a new methodology for anycast geolocation and enumeration (described in D4.3), with results that go beyond the industry standard (e.g., MaxMind). The technique not only enables new opportunities for Cloud troubleshooting (e.g., which anycast replica is the contacted one?), but also opens the way for the detection of violation of EU privacy rules (e.g., when private data is "cached" over Cloud nodes that are beyond EU boundaries) or BGP hijacking detection (e.g., when a previously unicast prefix is suddenly detected as anycast).

Since useful services such as Spamhaus were victims of unprecedented attacks recently, topping to a 300Gbps DDoS in March 2013, it follows that our *distributed* techniques are less easily vulnerable to such attacks, which is expected to become an important differentiation in the years to come.

7.1.8 NEC

NEC is actively discussing with business units, such as NetCracker, for the commercialization of BlockMon, the high-availability data analytics platform that has been developed throughout the project, towards tier-1 network operators in Europe. A meeting with a tier-1 network operator has already taken place, and a trail is ongoing.

7.1.9 TID

In telefonica, we are applying the techniques and methods developed during the mPlane project to optimize video delivery service and to detect QoS in our network. We are already in touch with product teams and we started applying our methods on real data.

Currently we have 1 patent filled: "Method and system to troubleshoot quality of service issues in mobile networks". In this patent we describe the core work that we have been working on regarding the "Mobile Network Performance Issue Cause Analysis" use-case.

7.1.10 FTW

Being FTW an applied research center, its main role is the transfer of research results and novel technological applications to the industry. FTW exploits the project results in two directions: firstly, by transferring them to its industrial partners, through industrial collaborative projects; secondly, by demonstrating the value of strategic research and therefore by strengthening FTW scientific excellence.

In terms of the industrial exploitation of mPlane results, FTW has already started to work with its

industrial partners with results directly obtained in the first two years of the project. In particular, we are currently applying the developed techniques and the obtained results on the automatic detection and diagnosis of network anomalies in an industrial project with two of our industrial partners, including a trans-national EU ISP (Telekom Austria) and a main EU vendor (NSN).

mPlane has also boosted the know-how and work done in FTW in the direction of big data analysis technologies for network monitoring and analysis. In particular, during the period we have pushed further the development of DBStream, and new industrial projects as well as new proposals for future research activities are in the short-time horizon.

In terms of strengthening FTW's scientific position, mPlane has resulted so far in a highly successful number of publications in international conferences, workshops, and journals. Indeed, FTW has already published 13 papers at major international conferences such as SIGCOMM, CoNEXT, ICDE, PAM, INFOCOM, etc., 4 papers at international workshops, and 1 magazine paper.

7.1.11 FHA

FHA will use the GLIMPSE system to do further research and use it in education (programming courses and the data networking course). The GLIMPSE results will be made available (in strictly anonymized form, if necessary as aggregated data) as open data. In addition, the GLIMPSE probe will be used to monitor the university network. The implementation experience will be fed back to the standardization community.

7.1.12 ULg

ULg has made all datasets collected freely available to the community.

Moreover, the extended version of scamper has been sent to CAIDA (who originally developed scamper). They can now use our ports in their own deployments/measurements.

7.1.13 ETH

ETH Zürich, as a university, sees exploitation of mPlane results primarily as a short-term goal, through the education of PhD and Master's students working on topics related to the project, especially in the development of passive and active measurement techniques for performance and connectivity measurement. However we see the key result and the long term impact of the project being primarily in the definition of a standards-ready protocol specification for the integration of heterogeneous measurement tools and data sources into a coherent measurement plane, the open-source release of a reference implementation of this protocol (led by ETH), as well as the release of high-quality, open-source software for specific measurements.

7.1.14 A-LBELL

A-LBELL exploitation of results is mostly directed towards feeding resource/traffic optimization processes with machine learning algorithm output. At a more foundation level, A-LBELL also works at identifying the limits of eager learning and determining how dynamic and distributed environ-

ments may take benefit of lazy learning methods.

7.1.15 FW

Fastweb is willing to adopt the mPlane system (probes, algorithms, repositories) in part of its network and exploit its troubleshooting capabilities.

References

- [1] *Who to Blame when YouTube is not Working? Detecting Anomalies in CDN Provisioned Services*, 2014.
- [2] A. Araldo and D. Rossi. Dissecting bufferbloat: Measurement and per-application breakdown of queueing delay. In *ACM CoNEXT'13, Student Workshop*, 2013.
- [3] A. Araldo and D. Rossi. A per-application account of bufferbloat: Causes and impact on users. In *5th International Workshop on TRaffic Analysis and Characterization (TRAC), Best Paper Award*, 2014.
- [4] A. Baer, P. Casas, L. Golab, and A. Finamore. Dbstream: an online aggregation, filtering and processing system for network traffic monitoring. In *TRAC, Nicosia, Cyprus, 08/2014* 2014. IEEE, IEEE.
- [5] A. Baer, A. Finamore, P. Casas, L. Golab, and M. Mellia. Large-scale network traffic monitoring with dbstream, a system for rolling big data analysis. In *International Conference on Big Data, IEEE BigData 2014*, Washington D.C., USA, 11/2014 2014. IEEE, IEEE.
- [6] Z. Ben-Houidi, G. Scavo, S. Ghamri-Doudane, A. Finamore, S. Traverso, and M. Mellia. Gold mining in a river of internet content traffic. In *6th International Workshop on Traffic Monitoring and Analysis, TMA*, London, 04/2014 2014. Springer, Springer.
- [7] I. Bermudez, S. Traverso, M. Mellia, and M. Munafò. A distributed architecture for the monitoring of clouds and cdns: Applications to amazon aws. *IEEE Transactions on Network and Service Management*, In press, 2014.
- [8] E. Bocchi, M. Mellia, and S. Sarni. Cloud storage service benchmarking: Methodologies and experimentations. In *3rd IEEE International Conference on Cloud Networking (IEEE CloudNet 2014)*, Luxembourg, 10/2014 2014. IEEE, IEEE.
- [9] P. Casas, P. Fiadino, and A. Bär. Understanding http traffic and cdn behavior from the eyes of a mobile isp. In *Passive and Active Measurements Conference, PAM 2014*, 2014.
- [10] P. Casas, P. Fiadino, A. Bär, A. D'Alconzo, A. Finamore, and M. Mellia. Youtube all around: Characterizing youtube from mobile and fixed-line network vantage points. In *EuCNC*, Bologna, IT, 06/2014 2014.
- [11] P. Casas, A. Sackl, S. Egger, and R. Schatz. Characterizing microsoft lync online in mobile networks: a quality of experience perspective. In *3rd IEEE International Conference on Cloud Networking*, Luxembourg, Luxembourg, 10/2014 2014. IEEE, IEEE.
- [12] P. Casoria, D. Rossi, J. Augé, M.-O. Buob, T. Friedman, and A. Pescape. Distributed active measurement of internet queuing delays. In *Passive and Active Measurement (PAM), Extended Abstract*, Los Angeles, USA, March 2014.
- [13] C. Chirichella, D. Rossi, C. Testa, T. Friedman, and A. Pescape. Passive bufferbloat measurement exploiting transport layer information. In *IEEE GLOBECOM*, 12/2013 2013.
- [14] L. Cittadini, S. Vissichio, and B. Donnet. On the quality of bgp route collectors for ibgp policy inference. In *IFIP Networking*, June 2014 2014.
- [15] S. Colabrese, D. Rossi, and M. Mellia. *Aggregation of Statistical Data from Passive Probes: Techniques and Best Practices*, volume 8406 of *Lecture Notes in Computer Science*, pages 38–50. Springer Berlin Heidelberg, 2014.
- [16] S. Colabrese, D. Rossi, and M. Mellia. Scalable accurate consolidation of passively measured statistical data. In *Passive and Active Measurement (PAM), Extended Abstract*, Los Angeles, USA, March 2014.
- [17] P. M. S. del Rio, D. Rossi, F. Gringoli, L. Nava, L. Salgarelli, and J. Aracil. Wire-speed statistical classification of network traffic on commodity hardware. In *ACM Internet Measurement Conference (IMC)*, 11/2012 2012.
- [18] M. Dell'Amico, D. Carra, M. Pastorelli, and P. Michiardi. Revisiting size-based scheduling with estimated job sizes. In *MASCOTS 2014, IEEE 22nd International Symposium on Modeling analysis and simulation of computer and telecommunication systems, September 9-11, 2014, Paris, France / Also published on ArXiv*, Paris, FRANCE, 03 2014.

- [19] G. Detal, B. Hesmans, O. Bonaventure, Y. Vanaubel, and B. Donnet. Revealing middlebox interference with tracebox. In *Proc. ACM/USENIX Internet Measurement Conference (IMC)*, October 2013.
- [20] D.Papadimitriou. Practical issues for the implementation of survivability and recovery techniques in optical networks. *Journal Optical Switching and Networking*, 14, 2014.
- [21] W. Du, Y. Liao, N. Tao, P. Geurts, X. Fu, and G. Leduc. Rating network paths for locality-aware overlay construction and routing. *IEEE/ACM Transactions on Networking*, 2014.
- [22] M. Dusi, R. Bifulco, F. Gringoli, and F. Schneider. Reactive logic in software-defined networking: Measuring flow-table requirements. In *5th International Workshop on TRaffic Analysis and Characterization (TRAC)*, Nicosia, Cyprus, 2014 2014.
- [23] P. Fiadino, A. D’Alconzo, A. Bär, A. Finamore, and P. Casas. On the detection of network traffic anomalies in content delivery network services. In *ITC26*, Karlskrona, Sweden, 09/2014 2014.
- [24] P. Fiadino, A. D’Alconzo, and P. Casas. Characterizing web services provisioning via cdns: The case of facebook. In *5th International Workshop on TRaffic Analysis and Characterization*, Nicosia, Cyprus, 08/2014 2014. IEEE, IEEE.
- [25] P. Fiadino, M. Schiavone, and P. Casas. Vivisecting whatsapp through large-scale measurements in mobile networks. In *SIGCOMM 2014*, Chicago, USA, 08/2014 2014. ACM, ACM.
- [26] A. Finamore, M. Mellia, Z. Gilani, K. Papagiannaki, V. Erramilli, and Y. Grunenberger. Is there a case for mobile phone content pre-staging? In *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies (Best Short Paper Award)*, New York, NY, USA, 12/2013 2013. ACM, ACM.
- [27] Y. Gong, D. Rossi, C. Testa, S. Valenti, and D. Taht. Fighting the bufferbloat: on the coexistence of AQM and low priority congestion control (extended version). *Elsevier Computer Networks*, 60, 2014.
- [28] L. Grimaudo, M. Mellia, E. Baralis, and R. Keralapura. Select: Self-learning classifier for internet traffic. *IEEE Transactions on Network and Service Management*, 11, 06/2014 2014.
- [29] C. Imbrenda, L. Muscariello, and D. Rossi. Analyzing cacheable traffic in isp access networks for micro cdn applications via content-centric networking. In *ACM SIGCOMM Information Centric Networks (ICN)*, Paris, FR, 09/2014 2014.
- [30] Y. Liao, D. Wei, P. Geurts, and G. Leduc. DMFSGD: A decentralized matrix factorization algorithm for network distance prediction. *IEEE/ACM Transactions on Networking*, 2013. to appear.
- [31] M. Luckie. Scamper: a scalable and extensible packet prober for active measurement of the Internet. In *Proc. ACM/USENIX Internet Measurement Conference (IMC)*, November 2010.
- [32] R. Mazloun, M.-O. Buob, J. Auge, B. Baynat, T. Friedman, and D. Rossi. Violation of interdomain routing assumptions. In *Passive and Active Measurement (PAM)*, Los Angeles, USA, 03/2014 2014.
- [33] G. Nassopoulos, D. Rossi, F. Gringoli, L. Nava, M. Dusi, and P. Santiago del Rio. *Flow Management at Multi-Gbps: Tradeoffs and Lessons Learned*, volume 8406 of *Lecture Notes in Computer Science*, pages 1–14. Springer Berlin Heidelberg, 2014.
- [34] D. Rossi, G. Pujol, X. Wang, and F. Mathieu. Peeking through the bittorrent seedbox hosting ecosystem. In *Traffic Monitoring and Analysis (TMA)*, 2014.
- [35] A. Rufini, M. Mellia, E. Tego, and F. Matera. Multilevel bandwidth measurements and capacity exploitation in gigabit passive optical networks. *IET Communications*, 2014.
- [36] A. Rufini, E. Tego, F. Matera, and M. Mellia. Bandwidth measurements and capacity exploitation in gigabit passive optical networks. In *Fotonica 2014*, 05/2014 2014.
- [37] E. Tego. Active measurements and limitations of tcp protocol during sla test. In *Poster session the Sixth Workshop on Trafβic and Monitoring Analysis (TMA)*, 04/2014 2014.
- [38] C. Testa and D. Rossi. Delay-based congestion control: Flow vs. bittorrent swarm perspectives. *Elsevier Computer Networks*, 60, 02/2014 2014.

- [39] B. Trammell. Evolving transport in the internet. *IEEE Internet Computing*, 09/2014 2014.
- [40] B. Trammell, P. Casas, D. Rossi, A. Bär, Z. Ben-Houidi, I. Leontiadis, T. Szemethy, and M. Mellia. mplane: an intelligent measurement plane for the internet. *IEEE Communications Magazine, Special Issue on Monitoring and Troubleshooting Multi-domain Networks using Measurement Federations*, 42, 05/2014 2014.
- [41] B. Trammell, D. Gugelmann, and N. Brownlee. Inline data integrity signals for passive measurement. In *Sixth International Workshop on Traffic Monitoring and Analysis (TMA 2014)*, 04/2014 2014.
- [42] S. Traverso, E. Tego, E. Kowallik, S. Raffaglio, A. Fregosi, M. Mellia, and F. Matera. Exploiting hybrid measurements for network troubleshooting. In *IEEE Networks*, Funchal, PT, 09/2014 2014. IEEE, IEEE.
- [43] Y. Vanaubel, J.-J. Pansiot, P. Mérindol, and B. Donnet. Network fingerprinting: Ttl-based router signatures. In *Proc. ACM/USENIX Internet Measurement Conference (IMC)*, October 2013.