



## mPlane

an Intelligent Measurement Plane for Future Network and Application Management

**ICT FP7-318627**

# Report on the mPlane “Training School”

<b>Author(s):</b>	POLITO	M. Mellia
	ENST	D.Rossi
	NEC	M. Dusi. M. Ahmed
	TID	I. Leontiadis. M. Varvello, L. Baltrunas
	FHA	R. Winter, M. Faath
	ULG	B. Donnet, K. Edeline

<b>Document Number:</b>	D7.4
<b>Revision:</b>	0.1
<b>Revision Date:</b>	31 Oct 2015
<b>Deliverable Type:</b>	Other
<b>Due Date of Delivery:</b>	31 Oct 2015
<b>Actual Date of Delivery:</b>	31 Oct 2015
<b>Nature of the Deliverable:</b>	(R)eport
<b>Dissemination Level:</b>	Public

**Abstract:**

The mPlane project - just like any other applied research project - depends on a community, on people outside the project to help establish the technologies built and developed as part of the project. A key to building a community is dissemination, education and training. The main dissemination outlet of the project is scientific publications. In terms of education and training, the mPlane consortium has made use of a number of activities to foster awareness of mPlane concepts and technologies and to train people in the use of these technologies. This document summarizes these activities.

**Keywords:** dissemination, training, education

## 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
1 Introduction.....	5
2 Training School Activities.....	6
2.1 1st PhD School on Big-Data Algorithms and Applications to Traffic Measurements (TMA2014).....	7
2.1.1 Response summary .....	8
2.2 4th Traffic Monitoring and Analysis summer school (TMA2014) .....	13
2.2.1 Response summary .....	14
2.3 Undergraduate Course on mPlane and GLIMPSE .....	17
3 Industrial Outreach.....	19
3.1 TMA'15 Industrial Worskhop.....	20
3.2 mPlane Final workshop at CoNext'15 .....	21
3.2.1 Guest speakers .....	21
3.2.2 Workshop Programme .....	22
4 Hands-on Training Activities.....	23
4.1 Coding an mPlane-Compliant Application/Proxy.....	24
4.2 Seminar on coding network monitoring applications .....	25
5 Conclusion.....	26

## 1 Introduction

The mPlane project has carried out a number of training events as part of the educational mission of the project. Apart from the usual dissemination activities, project partners have been active in three more or less distinct categories of training events: training schools, industrial outreach and hands-on coding tutorials. Each of these event categories is equally important as they achieve different things. The training school activities are usually a good way to create awareness in academic circles, to disseminate higher-level project ideas or focus on narrow but research-oriented aspects of a project. These activities help establishing project ideas in the scientific community. Then there are industrial outreach activities which are also very important, because at some point the mPlane technology needs to make the transition from research to deployed technology. Clearly, the industrial project partners cannot achieve wide adoption/deployment alone. Industrial outreach is about bringing the value of mPlane closer to the ultimate “customer”. Finally, there are hands-on coding tutorials, which have two important aspects. On the one hand, it gives people the opportunity to play with the technology on real systems, giving interested parties a better feeling for and training in mPlane technology. On the other hand, the interaction with mPlane “newbies” allows the developers inside the mPlane consortium to refine documentation, APIs and code based on the experience and feedback of the tutorial participants.

Overall, the mPlane project has done a number of such events which are all recorded in this deliverable. Except for one industrial outreach activity, all of these events have already taken place at the time of writing this document. The following is the list of these events:

- Training schools
  - 1st PhD School on Big-Data Algorithms and Applications to Traffic Measurements (TMA2014)
  - 4th Traffic Monitoring and Analysis summer school (TMA2014)
  - Undergraduate Course on mPlane and GLIMPSE
- Industrial outreach
  - TMA’15 Industrial Workshop
  - mPlane Final workshop at CoNext’15
- Hands-on coding tutorials
  - Seminar on how to code an mPlane-compliant application/proxy
  - Seminar on coding network monitoring applications

## 2 Training School Activities

The training school events organized by the mPlane consortium were naturally all carried out by academic partners. Overall three events were organized (in one case co-organized together with the FP7-BigFoot project) - two PhD schools and an undergraduate course. For the PhD schools, feedback forms were prepared and the evaluation of these are part of this deliverable. The undergrad course was an implementation-oriented course spanning a whole semester.

## 2.1 1st PhD School on Big-Data Algorithms and Applications to Traffic Measurements (TMA2014)

The school has been jointly organized by the FP7-mPlane and FP7-BigFoot European projects. The goal of the school was to offer participants insights into the technologies that are commonly referred to as “BigData”. The school’s homepage is still hosted on the mPlane website, see <https://www.ict-mplane.eu/public/phd-school-mapreduce>.

Topics addressed during the PhD school:

- The MapReduce programming paradigm
- Hadoop internals and libraries for BigData analysis
- Pig and other high-level programming languages for BigData
- Applications to Internet traffic measurement and analysis

In particular, considering the last point, the school focused on applications of Big Data technologies on the processing and extraction of information from passive measurements collected by Tstat, a well-known passive analyzer that has been developed in the context of mPlane.

The school was organized by POLITO on the Politecnico di Torino campus, during Wednesday 18th to Friday 20th of September 2013. The schedule covered three full days, organized with both lectures and lab session. The main presenters were Prof. Pietro Michiardi, from Eurecom and WP3 coordinator, and Dr. Antonio Barbuzzi from Eurecom and working on WP3 too. See <http://www.ict-mplane.eu/public/speakers> for their online CV.

The detailed program, available online at <http://www.ict-mplane.eu/public/program>, covered aspects of Big Data technologies, Hadoop Internals, and high-level languages, such as Pig and Pig latin, with hands-on session in laboratories during each afternoon. All material used during the class was and is made available on the mPlane website at <https://www.ict-mplane.eu/public/course-material>

We were targeting about 25 students, but we had to close the registration process in advance due to the overwhelming number of requests. At the end, 44 participants were accepted, of which about 50% were coming from mPlane partners, and 50% from Universities, Research Centers or companies in Europe. The complete list can be found at <http://www.ict-mplane.eu/public/school-participants>.

One year after the event, we collected feedback from the participants with Google forms. The idea was to assess if the methodologies and tools that have been presented during the school were found to be useful by the researchers. A set of questions were asked, of which some were used for the evaluation of other PhD Schools organized by mPlane. Each question allowed answers on a 1-5 scale. In total, 21 attendees replied, which we consider a good number of feedbacks to evaluate the success of the school. We first report a very quick overview of the percentage of answers getting the highest scores (4-5), and defer a detailed per-question breakdown to a later section (summary of Google forms).

1. How useful was the information presented during the school? Was this event better than what you expected, worse than what you expected, or about what you expected?

**(Answer: over 81% of the participants found the school to match their expectations or even exceed them)**

2. Considering the organization of the School, was the organization of this event well prepared?  
**(Answer: over 85% found the event quite to extremely well organized)**
3. Considering the topics that were presented during the school, how many of them did you find interesting and useful?  
**(Answer: over 95% found most or all the topics to be interesting and useful)**
4. How useful was the material presented at this event?  
**(Answer: over 81% found the material to be quite to extremely useful)**
5. How much of the training did you feel was hands-on?  
**(Answer: over 67% found that most or all the training contained hand-on parts)**
6. Considering the topics of the schools, how much have your skills improved because of the training at the event?  
**(Answer: over 37% found their skills improved significantly and 47% felt moderate improvement)**
7. Did you already have the chance to make use of learned material in practice?  
**(Answer: over 53% mentioned they use the tools and material learned in the school)**
8. How likely will you be able to make use of learned material in practice?  
**(Answer: 55% said they would likely or very likely use the material in practice)**
9. How would you rate the training in terms of level of expertise required to follow the school?  
**(Answer: 100% said they found the event was for a skilled to advanced audience)**

Overall, this represents a very positive feedback of the school. Some free text comments were also entered, and we report them as they are:

- The difficulty I had in following the training was that it assumed prior knowledge of Java and Eclipse environment. Of course it can be assumed one is able to code. Would have been great to receive a pointer to online training material on prerequisite for the course.
- Michiardi is a really great speaker!! I felt really impressed by him!

All answers are reported next, as the Google Form original printout.

### 2.1.1 Response summary

Analytics are publicly available at <https://docs.google.com/forms/d/1YniimbbW7nPSr8r2Q7LPqsVUiLtttpx/viewanalytics>. Here we report a detailed per-answer summary.




# 21 responses

[View all responses](#)


## Summary

**How useful was the information presented during the school? Was this event better than what you expected, worse than what you expected, or about what you expected?**



Better than expected	<b>11</b>	52.4%
Exactly as expected	<b>6</b>	28.6%
Slightly less than expected	<b>4</b>	19%
Significantly less than expected	<b>0</b>	0%
Failed my expectations	<b>0</b>	0%


**Considering the organisation of the School, was the organisation of this event well prepared?**



Extremely organized	<b>5</b>	23.8%
Quite organized	<b>15</b>	71.4%
Moderately organized	<b>1</b>	4.8%
Somewhat organized	<b>0</b>	0%


Not organized **0** 0%

**Considering the topics that were presented during the school, how many of them did you find interesting and useful?**




All	<b>6</b>	28.6%
Most	<b>14</b>	66.7%
About half	<b>0</b>	0%
A few	<b>1</b>	4.8%
None	<b>0</b>	0%

**How useful was the material presented at this event?**




Extremely useful	<b>4</b>	19%
Quite useful	<b>13</b>	61.9%
Moderately useful	<b>4</b>	19%
Somewhat useful	<b>0</b>	0%
Not useful	<b>0</b>	0%

**How much of the training did you feel was hands-on?**




All of it	<b>4</b>	19%
Most of it	<b>10</b>	47.6%
About half	<b>6</b>	28.6%
Some of it	<b>1</b>	4.8%
None	<b>0</b>	0%

**Considering the topics of the schools, how much have your skills improved because of the training at the event?**




Significantly	<b>3</b>	14.3%
A lot	<b>5</b>	23.8%
Moderately	<b>10</b>	47.6%
A little	<b>3</b>	14.3%
Not at all	<b>0</b>	0%

**Did you already had the chance to make use of learned material in practice?**



Yes, I use them very often	<b>1</b>	4.8%
Sometimes I had the chance to use them	<b>10</b>	47.6%
Not yet, but I hope to have the chance to use them in the near future	<b>6</b>	28.6%
I hardly make use of them	<b>4</b>	19%


**How likely will you be able to make use of learned material in practice?**



Very likely	<b>3</b>	15%
-------------	----------	-----

Likely	<b>8</b>	40%
Moderately likely	<b>6</b>	30%
Not likely	<b>3</b>	15%

**How would you rate the training in terms of level of expertise required to follow the school?**




Advanced	<b>4</b>	20%
Skilled	<b>16</b>	80%
Unskilled	<b>0</b>	0%

**Additional feedback/criticism/suggetions**

The difficulty I had in following the training was that it assumed prior knowledge of Java and Eclipse environment. Of course it can be assumed one is able to code. Would have been great to receive a pointer to online training material on prerequisite for the course. Michiardi is a really great speaker!! I felt really impressed by him!

**Number of daily responses**



## 2.2 4th Traffic Monitoring and Analysis summer school (TMA2014)

### Date and Place:

- April 15th-16th, London UK

### Chair:

- Dario Rossi (ENST)

### Speakers:

- Benoit Donnet and Korian Edeline (ULG)  
Active inference: Revealing Middleboxes with Tracebox
- Alessandro Finamore (POLITO)  
Passive inference: Troubleshooting the Cloud with Tstat

The 4th PhD school on Traffic Monitoring and Analysis (TMA), was organized by mPlane partners however without mPlane funding support. The school was kindly sponsored by ACM SIGCOMM with about 25k USD that were mostly used to offer travel support to PhD students attending the school. The PhD school took place right after the TMA conference and school attendees had the chance to follow the TMA workshop at no cost. Students interacted with TMA participants through a poster sessions, which gave them a chance to present their ongoing work to TMA attendees and receive valuable feedback.

Methodology-wise, the school continued 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/>). School talks were followed by hands on laboratory sessions where PhD students had the chance to put cutting-edge methodologies they just were exposed to into practice – which both reinforces the learning process and drives the cost to start new research work in the TMA domain down.

Content-wise, the 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. We only report here an itemized set of topics here, since the detailed program has been reported in a previous deliverable and is available along with tools and datasets at <http://networks.eecs.qmul.ac.uk/news/tma-2014/phdschool/>

- Active inference: revealing middleboxes with tracebox
- Achieving scale: Large scale active measurement from PlanetLab
- Passive inference: Troubleshooting the Cloud with Tstat
- Close to users: Performance measurement via Fathom

After the event, we collected feedback from the participants with Google forms, with about 75% recall (31 out of 44 participants). The same set of questions that were asked for other events were used. Each question allowed answers on a 1-5 scale. We first report a very quick overview of the percentage of answers getting the highest scores (4-5), and defer a detailed per-question breakdown to a later section (summary of Google forms).

1. How useful was the information presented at TMA school ? Was this event better than what you expected, worse than what you expected, or about what you expected?  
**(Answer: over 75% of the participants found the school exactly as or better than their expectations)**
2. How useful was the information presented at TMA school ?  
**(Answer: over 80% found the event quite to extremely useful)**
3. How much of the training did you feel was hands-on?  
**(Answer: over 60% found most or all the event was hands-on)**
4. How organized was the information presented at TMA school ?  
**(Answer: over 65% found the event to be quite to extremely-well organized)**
5. How many of the objectives of the TMA school were met?  
**(Answer: over 80% found that most or all the objectives to be met)**
6. How much have your skills improved because of the training at the TMA school?  
**(Answer: over 20% found their skills improved significantly and 50% felt moderate improvement)**
7. How likely will you be able to make use of learned material in practice?  
**(Answer: over 36% mentioned they would likely or very likely use tools learned in the school)**
8. Will you mention/recommend TMA related technologies to colleagues after this event?  
**(Answer: 90% said they would likely or very likely recommend the TMA technologies)**
9. How would you rate the training in terms of level of expertise required?  
**(Answer: over 95% said they found the event was targeting a skilled to advanced audience)**

Overall, a very positive feedback of the school. More details are reported next.

### 2.2.1 Response summary


Analytics are publicly available at [https://docs.google.com/forms/d/1PJ\\_nSQTjAsoI0Ch4dDnF2uqERwgCI1G9viewanalytics](https://docs.google.com/forms/d/1PJ_nSQTjAsoI0Ch4dDnF2uqERwgCI1G9viewanalytics). Here we report a detailed per-answer summary.

# 31 responses

[View all responses](#)


## Summary

How useful was the information presented at TMA school ? Was this event better than what you expected, worse than what you expected, or about what you expected?




Better than expected	8	25.8%
Exactly as expected	16	51.6%
Slightly less than expected	7	22.6%
Significantly less than expected	0	0%
Failed my expectations	0	0%

How useful was the information presented at TMA school ?




Extremely useful	6	19.4%
Quite useful	19	61.3%
Moderately useful	4	12.9%
Somewhat useful	2	6.5%
Not useful	0	0%

How much of the training did you feel was hands-on?




All of it	1	3.2%
Most of it	18	58.1%
About half	10	32.3%
Some of it	2	6.5%
None	0	0%

How organized was the information presented at TMA school ?




Extremely organized	8	25.8%
Quite organized	12	38.7%
Moderately organized	8	25.8%
Somewhat organized	3	9.7%
Not organized	0	0%

How many of the objectives of the TMA school were met?




All	5	16.7%
Most	19	63.3%
About half	6	20%
A few	0	0%
None	0	0%

**How much have your skills improved because of the training at the TMA school?**




Significantly	2	6.5%
A lot	5	16.1%
Moderately	15	48.4%
A little	9	29%
Not at all	0	0%

**How likely will you be able to make use of learned material in practice?**




Very likely	4	12.9%
Likely	9	29%
Moderately likely	17	54.8%
Not likely	1	3.2%

**How likely will you be able to make use of learned material in practice?**




Very likely	3	10%
Likely	8	26.7%
Moderately likely	17	56.7%
Not likely	2	6.7%

**Will you mention/recommend TMA related technologies to colleagues after this event?**



Very likely	13	43.3%
Likely	14	46.7%
Moderately likely	3	10%
Not likely	0	0%

**How would you rate the training in terms of level of expertise required?**



Advanced	10	32.3%
Skilled	20	64.5%
Unskilled	1	3.2%

**Additional feedback/criticism/suggestions**

Thanks for the good organisation. However, I expected at least the same amount of hands-on activities as in the previous year. Maybe it is a problem with my expectations. A suggestion is to make the attendees to work together and generate/create something to be presented in the end of each day.

I

I suggest to deal Network Security field in the next event..

it is a five star event :)

I think the pace of the program could have been slower

more hand on experience with previously provided theory would be appreciated, more dynamic presentations with less time would be also



## 2.3 Undergraduate Course on mPlane and GLIMPSE

The Computer Science curriculum of the University of Applied Sciences Augsburg (FHA) contains two group projects in the 4th and 6th semester, where students in teams of 8 work on a well-defined, design and implementation-focused, semester-long project. mPlane-related work, in particular the work contained in work package 2 (Programmable Probes) lent itself to this kind of project.

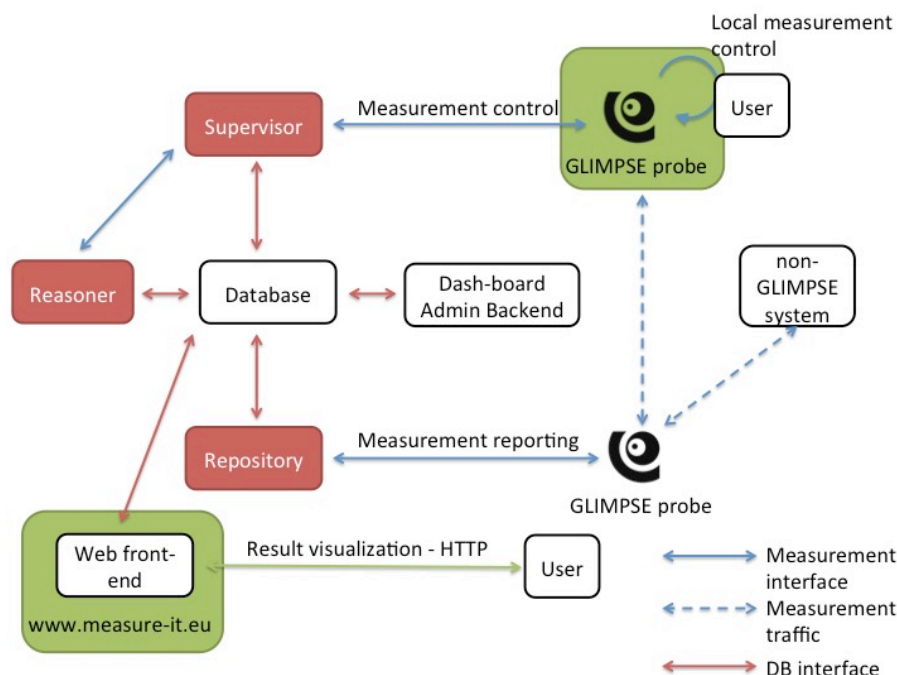


Figure 2.1: GLIMPSE overview

The overall GLIMSPE system—an implementation of the mPlane architecture developed at FHA—was too large and complex to have as a project. The probe and the web front-end were chosen as primary subjects for development work as highlighted (green areas) in figure 2.1. In addition, since the GLIMPSE probe is a C++(with Qt) cross-platform application, setting up a build system was also part of the project as well as project coordination amongst the students. Ideally, these projects have a well-defined task, where the end result is delivered at the end of the semester (including documentation). In order to achieve a reasonable outcome, frequent (weekly) meetings were scheduled with all team members.

The intent of the project, developed during the Academic Year 2015-2016, was twofold. On the one hand, the students of course should learn how to work as a group, should acquire skills on both a technical and a social level. This is the intended outcome from the educational side. On the other hand, the project was a kind of test balloon for the future of the GLIMPSE project. All software within the project will be open sourced. The software should not just be archived but should live on. Being actively maintained and further developed is the key part in this. The project helped to estimate how difficult it is to have new developers come on board, how long it takes for them to become productive, which aspects of the project need better documentation and how long it takes to finish certain tasks (like adding a new measurement to the system).

During the project new measurements were implemented (not all of which are part of today’s ver-

sion of GLIMPSE), the build system was set up (which is still being used to date), a result visualization library was implemented which ultimately was not used but gave valuable insights into the cross-platform aspects of result visualization on the graphical probes. This project had in the end little influence on the active GLIMPSE code base (which was not a primary goal) but was a great learning opportunity for the students. In particular:

- students had to work with an existing code base which had to be understood
- students had to integrate code into this code base, adhering to the project’s coding guidelines
- students had to collaboratively work within a given CI environment
- students got an insight view into current research questions and network measurement tools and technologies
- the project participants had to present their results at a booth in a student project fair and had to explain their work and the mPlane architecture implemented by GLIMPSE to fellow students and professors

On a final note, from the experience gained during the project an additional optional-to-pick course was designed and is now part of the computer science curriculum. This course is on cross-platform development using C++ and Qt. It turned out, that this particular combination is a great learning base as it touches upon many different things. An additional hope is of course, that there will be students that can (in the future) do work on the GLIMPSE code base as part of a thesis, semester project or otherwise.

### 3 Industrial Outreach

At a certain stage in a research project, reaching out to key stakeholders is important to create awareness of developed technology, to receive feedback on system design and to allow for adjustments and fine-tuning based on these stakeholders’ interests.

These activities in the case of mPlane are industrial outreach activities to showcase mPlane technology. One such event already took place recently (in April 2015), which was an industrial workshop co-located with TMA. A second event is planned for the 30<sup>th</sup> of November. mPlane’s final workshop will be held co-located with the CoNEXT conference, introducing mPlane to a much wider audience.

## 3.1 TMA'15 Industrial Worskhop

mPlane’s Industrial workshp took place on April 22, 2015, in conjuction with TMA’15.

The intent of the workshop was to bring together invited speakers, academic researchers and industrial partners in order to talk about practical use cases relevant to operators and providers. This includes the discussion of topics about proactive and reactive troubleshooting for network communications at scale. Particular attention was dedicated to discussing the definition of a measurement plane for the Internet from a regulatory perspective, to shed light on Internet quality and guarantee net neutrality.

There were 28 participants, about 40% from mPlane, 40% from the industry and 20% from academia.

The program of the conference contained two keynote talks:

- J. Cesar, CTO of Zhilabs on Automatic Root Cause Analysis in Convergent networks
- M. Molina on Cooperative inter-operator traffic measurement frameworks: technical challenges and barriers for industrial adoption.

In addition to the keynotes, there were a number of industrial participants who showcased their technologies (Zhilabs, Talaia Solutions, Radiopt, Callstats.io). Furthermore, mPlane partners presented the mPlane architecture and how industrial partners can benefit from it. A live demo of the mPlane Protocol and Reference Implementation was shown, including a simple setup where active probes are orchestrated by a Reasoner with the goal of automatically identifying changes in congestion on backbone links.

There was also a lively panel discussion that included D. Papagianaki, C. Diot, B. Trammel, M. Molina, J. Cesar. The main topics discussed here were i) Ethics, privacy and security, ii) current and future measurement methodologies and iii) cooperation and incentives to join a measurement plane. Finally, there was a poster session where the main use cases of mPlane were explained to the participants and feedback was given by our industrial and academic audience.

The full program can be found here: <http://tma-2015.cba.upc.edu/mplane-program>

## 3.2 mPlane Final workshop at CoNext'15

The mPlane final workshop has been organized to coincide with the CoNEXT<sup>1</sup> conference which is being organized by NEC and will be taking place in Heidelberg during the week of 1-4 Dec. 2015.

The goal of the mPlane final workshop is to introduce the challenges of Internet measurement to a wider audience, highlight the contributions that mPlane has made towards this goal, and disseminate the software produced. As such, mPlane partners will give talks concerning the various components of the mPlane architecture and demonstrate these components and use cases that are built on them to the audience.

### 3.2.1 Guest speakers

Aside from the mPlane partners, the workshop will include three talks from prominent researchers in the fields of network measurement. In order of appearance, the following talks will be presented by the invited speakers.

**Balachander Krishnamurthy**, AT&T Labs: *Is measurement still an afterthought?*

**Abstract:** In the early days of ARPAnet it was possible to trace passage of a single packet through the network. Traffic between nodes could be directly measured and routers could be queried for current statistics. But only in the late 1980s were SNMP and traceroute created (and are still in use!). With fragmentation and decentralization, measurement became effectively a second-class citizen. Thus, architectural innovations come without measurement as a fundamental concept requiring reinvention of tools and techniques. Most of the work is now guided by needs of applications. I believe that we should identify measurement primitives that can aid in quicker construction of relevant tools in new application areas or novel architectures. This talk is aimed at pushing the community towards such a search using one application area as exemplar: SDN.

**Fabian E. Bustamante**, Northwestern University: *Content distribution on next generation cellular networks*

**Abstract:** Smartdevices are becoming the primary, and for many only, Internet point of access for an ever-larger fraction of users. Nearly a quarter of web traffic is mobile, and content consumption on mobile devices is expected to grow fourfold by 2018. Much of this content is served by content delivery networks (CDNs) that, as in the rest of the Internet, replicate content in servers around the world and redirect clients to nearby replicas. Replica selection is typically done following some well-established heuristic, such as client IP or the proximity between most hosts and their local DNS servers. Unlike in the rest of the Internet, however, these heuristics are generally ineffective, with their ineffectiveness becoming obvious in fast, next generation cellular networks. In this talk, I will discuss our work on understanding next generation cellular network infrastructure and configuration and evaluating its impact on content distribution and users's quality of experience. I will then sketch a new approach to replica selection that acknowledges the centrality of cellular gateways in cellular network routing and present early results from a realization of this approach as a client/CDN cooperation, i.e. without the cooperation of cellular operators.

**Alberto Dainotti**, Center for Applied Internet Data Analysis (CAIDA): *BGPStream, A framework for the historical analysis and real-time monitoring of BGP data*

**Abstract:** I will present the design and implementation of BGPStream, an open-source software framework (available at [bgpstream.caida.org](http://bgpstream.caida.org)) for the analysis of historical and live Border Gate-

<sup>1</sup><http://conferences2.sigcomm.org/co-next/2015/#!/home>

way Protocol (BGP) measurement data. Although BGP is a crucial operational component of the Internet infrastructure, and is the subject of fundamental research (in the areas of performance, security, topology, protocols, economy, ...), there is no standard and easy way of processing large amounts of BGP measurement data. BGPStream fills this gap by making available a set of API and tools for processing large amounts of live and historical data thus supporting investigation of specific events, rapid prototyping, and building complex tools and efficient large-scale monitoring applications (e.g., detection of connectivity disruptions or BGP hijacking attacks). I will describe the design choices and challenges in the development of BGPStream. I will present how the components of the framework can be used in different applicative scenarios, and I will describe the development and deployment of complex services for global Internet monitoring that we built on top of it.

### 3.2.2 Workshop Programme

The final workshop is organized to handle approximately 30 participants, as well as the partners and speakers. Details can also be found on the workshop page on <https://www.ict-mplane.eu/public/mplane-workshop-platform-measuring-internet> and are summarized below.

- 9:00am Registration and coffee
- 9:15am Overview of the mPlane project [M. Mellia, project coordinator]
- 9:45am mPlane architecture and principles (including reference open implementation) [B. Trammell, leader of mPlane architecture work package]
- 10:15am Keynote by Balachander Krishnamurthy, AT&T Labs - Research “Is measurement still an afterthought?”
- 11:00 Break
- 11:20am mPlane live demo [F. Invernizzi, leader of mPlane deployment work package]
- 11:45am Keynote by Fabián E. Bustamante, Northwestern University “Content distribution on next generation cellular networks”
- 12:30pm Lunch
- 1:30pm mPlane components: probes [A. Bakay, leader of mPlane work package on programmable probes]
- 2:00pm mPlane components: repositories [P. Michiardi, leader of mPlane work package on large scale data analysis]
- 2:30pm mPlane components: reasoners [P. Casas, leader of mPlane work package on iterative and adaptive analysis]
- 3:00pm Break
- 3:15pm Keynote by Alberto Dainotti, Center for Applied Internet Data Analysis (CAIDA) “BGP-Stream, A framework for the historical analysis and real-time monitoring of BGP data”
- 4:00pm Live demos of mPlane use cases
  - Estimating Content and Service Popularity
  - Active Measurements for Multimedia Content Delivery
  - Quality Experience for Web Browsing
  - Mobile Network Performance Issue Analysis
  - Anomaly detection and root cause analysis in large-scale networks
  - Verification of SLA
  - Passive Content Curation
- 5:30pm Wrap-up and final remarks

## 4 Hands-on Training Activities

The PhD schools had practical exercises in them but more intense hands-on activities were organized by the project. These were very focussed and in the context of the METRICS Initial Training Networks (ITN). The main goal of the METRICS project is to study the factors that influence the understanding of the performance of the global Internet, which covers the scope of mPlane.

From Monday August 25<sup>th</sup>, 2014 to September 5<sup>th</sup>, 2014 at Université catholique de Louvain in Louvain-la-Neuve, Belgium, was held the METRICS Bootcamp 2014. mPlaners were invited to give tutorials. The objective of the METRICS Bootcamp was to introduce new PhD students to the broad research domain of Internet measurements. The goal of this event was to give students an overview of real work on Internet measurements, training on communication skills (writing scientific papers, giving oral presentations and presenting results as graphs). This Bootcamp was targeted at project members but was open to external experts who also work on Internet measurements. The scientific program was a mix between real work on Internet measurements and training on communication skills.

## 4.1 Coding an mPlane-Compliant Application/Proxy

In parallel to the Blockmon seminar during the Initial Training Networks (ITN) METRICS project (see Sec. 4.2), we proposed a tutorial on how to code an mPlane-compliant application/proxy. This tutorial was given by Mr. Korian Edeline (ULG).

First, Mr. Edeline introduced the mPlane project, the objectives pursued by the project (and the kind of problems mPlane is trying to solve) and the basic mPlane architecture (interactions between the probe(s), the repository, and the intelligent reasoner).

Secondly, Mr. Edeline went deeper into the mPlane architecture, focusing on the fact that everything, in mPlane, is a “component”. The main components (probes, repositories, supervisor) were described.

Thirdly, Mr. Edeline explained, based on the mPlane architecture, how anyone can write a mPlane-compliant proxy. As an example, together with the audience, a set of CISCO ping agents were built.



## 4.2 Seminar on coding network monitoring applications

During the seminar on “Blockmon: Architecture and how to code network monitoring applications”, M. Dusi from NEC presented the architecture of Blockmon, one of the probes for collecting Internet measurements developed under mPlane and available as open source. During the seminar, the people attending had the opportunity to learn about how to code network monitoring applications with Blockmon, with a focus on designing for performance. The architecture of Blockmon, based on building blocks that can be combined to achieve complex monitoring applications, turned out to be particularly suited for this purpose.

After the detailed presentation, the students could also join a hands-on experience by coding network monitoring applications themselves through the Blockmon platform. Some students expressed their satisfaction with the seminar and were determined to start using Blockmon for their research activities in the field of network monitoring and measurements.

## 5 Conclusion

The mPlane project organized quite a number of training/dissemination/community-building activities, which were described in this document. All of these have received favorable reviews by the participants and a large number of the questionnaire answers indicated that taught material is being used today. There is only a single event left before the project finishes which is the final workshop on the 30<sup>th</sup> of November in Heidelberg, which will contain a large number of mPlane demonstrations.

## References