mPlane

an Intelligent Measurement Plane for Future Network and Application Management

ICT FP7-318627

Demonstration Records

Author(s):	Author names
POLITO	S. Traverso
FUB	E. Tego, F. Matera
TI	F. Invernizzi, O. Jabr
ALBLF	Z. Ben Houidi
EURECOM	M. Milanesio
ENST	D. Rossi
NEC	M. Ahmed
TID	I. Leontiadis, M. Varvello
NETvisor	B. Szabo
FTW	P. Casas
FHA	M. Faath, R. Winter
ETH	B. Trammell
FW (ed.)	A. Fregosi, A. E. Kahveci, E. Kowallik, G. P. Mattellini, C. Meregalli, S. Raffaglio A. Sannino, M. Scarpino

Document Number:	D6.3
Revision:	2.0
Revision Date:	15 January 2016
Deliverable Type:	DEM
Due Date of Delivery:	31 December 2015
Actual Date of Delivery:	10 May 2016
Nature of the Deliverable:	(R)eport
Dissemination Level:	Public

Abstract:

This public deliverable provides a summary of all the demonstration events in which mPlane use cases and software systems were demonstrated, outlining the effects of those demonstrations on the respective audiences, in particular the feedback received from them. Overall, this report contains details on seven events that the project organized or participated in. A special focus of the report is on the final mPlane demonstration event, where all use cases of the project where shown plus a number of additional demonstrators that where built inside the project. Additional information on these events can be found at a number of references throughout the text where e.g. video material or other information can be found.

Contents

1 Introduction

1.1 Purpose of the document

This deliverable provides the demonstration reports of the events & conferences where the capabilities and outcomes of the mPlane project were presented. It describes the results of the mPlane demonstration activities, how the demonstrations have been carried out and what the audience feedback was.

1.2 Structure of the document

This deliverable describes the demonstration of the mPlane project, its software and its use cases at selected conferences. The deliverable is organized into four main sections, which are:

- **Introduction**
- **Context of the respective demonstrations**: this section provides the scope of the demonstrations and a list of demonstration events.
- **Demonstration event reports**: this section is divided into two subsections, namely a report on early demonstration events and a more detailed report of the final demonstration event at the end of the project. The former describes smaller demonstration activities carried out during the project period, sometimes early in the project. The latter summarizes the final demonstration event, where all use cases were demonstrated using the fully functional mPlane protocol.
- **Conclusion**

2 Context of the respective demonstrations

The wide variety of applications and the sheer complexity of the Internet's distributed network lead to operational frailty and results in severe difficulties when trying to identify and track the root cause of performance and availability issues.

In this context, to elucidate the current and obscure dynamics of the Internet, the mPlane project has focused on building an Internet measurement plane, by designing a flexible, open platform that allows one to collect, store and process measurement data collected from the Internet. Through the demonstration events described in this document, the project presented its outcomes in an applied and easily graspable manner, which aimed to provide solutions to identifying the root cause of the issues mentioned above from both from a provider and from a user point of view. The aim in particular was to reach out beyond the scientific community with events such as industrial workshops.

2.1 List of demonstration events

The capabilities and outcomes of the mPlane project were presented at various events. The main events were:

- mPlane industrial workshop at TMA (Barcelona, April 22, 2015)
- mPlane at Future X days (Nozay, June 11-12, 2015)
- mPlane at EuCNC (Paris, June 29-July 2, 2015)
- mPlane at Journée du Conseil Scientifique de l'Afnic (Paris, July 09, 2015)
- mPlane at the Innovation Workshop at Telecom Italia Labs (Turin, July 14, 2015)
- mPlane final workshop (Heidelberg, November 30, 2015)
- 9th EAI International Conference on Performance Evaluation Methodologies and Tools (Berlin, December 14-16, 2015)

3 Demonstration event reports

3.1 Report on early and smaller demonstration events

3.1.1 mPlane industrial workshop at TMA

mPlane has been presented at the mPlane industrial workshop co-located at the 7th International Workshop on Traffic Monitoring and Analysis (TMA 2015), Barcelona, April 22, 2015.

During the mPlane industrial workshop, a demo session took place. Its main aim was to demonstrate the capabilities of the mPlane system in a live environment, showing multiple probes (Telecom Italia DATI, generic pinger and HTTP latency probe, PoliTO Tstat) interacting with a supervisor, managed by means of different interfaces (python CLI, nodejs CLI, WEB GUI).

The capabilities of a reasoner developed during integration activities were presented. As part of the demonstration, network impairments were created that the reasoner discovered and indentified by analyzing measurements of passive probes (pingers). All communications between the involved elements (supervisor, probes, reasoner, GUI) was based on the mPlane protocol running over HTTPS.

More information on this event is available at http://tma-2015.cba.upc.edu/mplane.

3.1.2 mPlane at Future X days

WeBrowse (representing the passive content curation use case) was demonstrated during the Bell Labs Future X days (Nozay, June 11-12, 2015). The FutureX days are open door days where Bell Labs researchers and their partners demo their work to visitors from the industry as well as to ALU business units.

The passive content curation use case aims at helping users discovering relevant content in the web from simple passive observation of network traffic. The prototype was running live at the campus network of Politecnico di Torino. Since the best way to test WeBrowse is to check how users react to its content, the FutureX days was a great opportunity to test how people reacted to WeBrowse's promoted web stories and articles.

Overall, the stories promoted by WeBrowse attracted the smiles and the curiosity of many visitors. Interesting feedback and discussions occurred with many visitors, particularly, around the issues of privacy, the need for or harm of personalization, and the deployment of WeBrowse in other networks and environments.

3.1.3 mPlane at the EuCNC

The mPlane project was one of the exhibitors at the European Conference on Networks and Communications (EuCNC 2015), in Paris, June 29-July 2, 2015. The presented live demos were:

- Demonstration of the mPlane architecture and protocol: the goal was to demonstrate the approach defined by the mPlane architecture that allows a seamless integration of probes,

repositories, and reasoner, orchestrated by a supervisor

- Demonstration of the system integration and reasoning (YouTube video)
- Demonstration of a novel methodology to identify and locate IP Anycast addresses using active measurements
- Demonstration of a novel content curation approach based on passive observation of traffic (WeBrowse)
- Demonstration of a novel YouTube QoE-based monitoring probe for smart phones
- Demonstration of over-the-top (OTT) video and mobile YouTube troubleshooting

Due to the limited space at the booth, it was impossible to demonstrate all the mPlane use cases, hence, before the event, an official mPlane YouTube channel was set up. It contains videos of the remaining demos, that were played in a loop on one screen and were made part of the event this way. An image gallery of the EuCNC and feedback after the event is available as a blog entry on the mPlane website.

YouTube channel: https://www.youtube.com/channel/UCHGS6UlUKvGZTyt5DemmPaw
Blog entry in mPlane website: https://goo.gl/PGlMpH

mPlane demos were generally well received; the crowd was pleased with our live presentations such as an on-emachine version of the Multimedia Content delivery UC and they also liked the YouTube channel (YoMo app and OTT demonstrations have scored 80 views each). Despite of not presenting the fully functional architecture, the prerecorded demo played was ready to balance this shortcoming. However, as a consequence, the demos of each use case were felt as "standalone" demos by the audience, and not part of a single architecture. We inferred that a demo of the mPlane architecture and reference implementation should have been systematically given before going into details of any use case (which was difficult at times when people were arriving when the demo had already started), which we did in following events.

Other suggestions were to have more interactive demos -- specifically, people suggested to let participant download and install apps on their devices to interact with the demo (e.g., YoMo app) and let the demo be more interactive by reacting to user actions (i.e. anycast, passive measurement). These recommendations were followed up to ameliorate mPlane software even after the end of the project (YoMo app is now in Google play store with over 100 downloads; the anycast iGreedy software integrates RIPE Atlas streaming APIs to analyze and display data as they are received, etc.)

It must be also admitted that due to booth space and time limitations, and also because of the sheer number of projects demonstrated around, most of the visitors were not able to grasp the essence of the mPlane concept from our short introduction and posters, rather focusing on individual elements not seeing it in one piece.

3.1.4 mPlane at Journée du Conseil Scientifique de l'Afnic

mPlane was invited to the #JCSA15, the 5th "Journée du Conseil Scientifique de l'Afnic" [2] which focused on Internet measurement. Afnic is the French "Association Française pour le Nommage Internet en Coopération", which handles the .fr top level domain name (TLD), with about 2.8 million recorded names, and it handles 17 new generc TLDs as well.

The #JCSA15 was the 5th in a series of annual events, that receives a very large coverage, with live streaming via webcast [4], and a lot of Twitter activity [5]. The event was structured in two parts, with a morning session dedicated to a RIPE Atlas tutorial, and an afternoon session with scientific seminars. Dario Rossi (ENST) represented mPlane [4], covering the essential aspects of the project, in a session where also Daniel Karrenberg (CSO de RIPE) presented RIPE Atlas [3].
Gathered feedback was mostly positive. Daniel Karrenberg, Chief Scientist of RIPE NCC, donated a significant amount of credits to mPlane, and invited the consortium to submit proposals to RIPE71 (finally we had 1 pleanary and 3 MET presentations) . No negative feedback has been received directly. Rather, informal feedback from the organizer said that the mPlane presentation was the one receiving the highest scores (feedback was collected during the event but is not open to the public).

3.1.5 mPlane at the Innovation Workshop at Telecom Italia Labs

PoliTO gave an invited talk at the Innovation Workshop at Telecom Italia Labs, July 14, 2015. The aim of the talk "WeBrowse: a Passive Content Curation System" was disseminating the idea of a crowdsourced content curation system based on the passive analysis of HTTP traffic flowing in real operational networks. During the talk, PoliTO presented the architecture of the system and the modules composing it. Lastly, as a proof of concept, the prototype was presented, namely WeBrowse (http://webrowse.polito.it), which is deployed in the premises of Politecnico di Torino and builds on the mPlane protocol to support the interactions among the several modules composing the system.
Overall, the feedback from the audience was very positive. All people in the room were rather happy, or surprised by the methodological approach and the content promoted by WeBrowse. In particular, Christian Cocozza and Maurizio Siviero (Telecom Italia Labs) asked for more information about the exploitation plans regarding WeBrowse, and about the possibility to deploy it in networks other than PoliTO.
Moreover, part of the audience raised some interest about the algorithms and mechanisms WeBrowse implements to protect users' privacy.

3.1.6 9th EAI International Conference on Performance Evaluation Methodologies and Tools

The GLIMPSE measurement probe was presented by FHA at an invited talk and demonstration at the 9th EAI International Conference on Performance Evaluation Methodologies and Tools in Berlin, December, 2015. The talk introduced the mPlane architecture, the capabilities of the developed probe as well as the benefits of using a distributed measurement platform to gain network insights and troubleshooting information. The demonstration took place in-between talks as poster sessions, where FHA showed the probe itself as well as multiple measurement results from an HTTP download speed measurement campaign.

The feedback received showed that it is necessary to improve the possibilities for end users of GLIMPSE to schedule own measurements. The demonstrated version only allowed users to schedule latency measurements with a simple periodic timing, whereas multiple participants of the conference asked for bandwidth and traceroute measurements with complex measurement schedules to monitor their own infrastructure. The current deployment of GLIMPSE allows FHA to schedule

these kind of measurements to test the probe, supervisor and repository. Future development has to focus on user-scheduled measurements to make the platform more attractive for end users.

More information on this event are available at http://valuetools.org/2015/show/home.

3.2 Final demonstration event report

The mPlane workshop took place at NEC Laboratories Europe on November 30, 2015, in conjunction with CoNEXT 2015, the 11th International Conference on Emerging Networking Experiments and Technologies.

The workshop featured presentations and talks from measurement experts who spoke about topics related to the main concepts of building an Internet measurement plane. These talks by external speakers were supplemented by presentations about the entirety of the mPlane reference architecture and components. Finally, all mPlane use cases were demonstrated as well as a set of additional demonstrations that were developed as part of the project.

During the workshop, each participant was given a set of brochures that included the agenda, an overview of the project, and detailed information about the architecture and use cases.
Following the presentation sessions, workshop attendees moved to the dedicated demonstration area, where the applicability and functionality of the use cases and the flexibility of the overall system was shown. In the demonstration area, the participants had the opportunity to discuss and learn more about the specific use cases, through the prepared posters and parallel live demos.

By including invited speakers from industry, this workshop brought together mPlane use cases relevant to operators and providers and potential users of the mPlane system. This gave opportunity to discuss topics about proactive and reactive troubleshooting for network communications at a scale.

3.2.1 Attendee profile

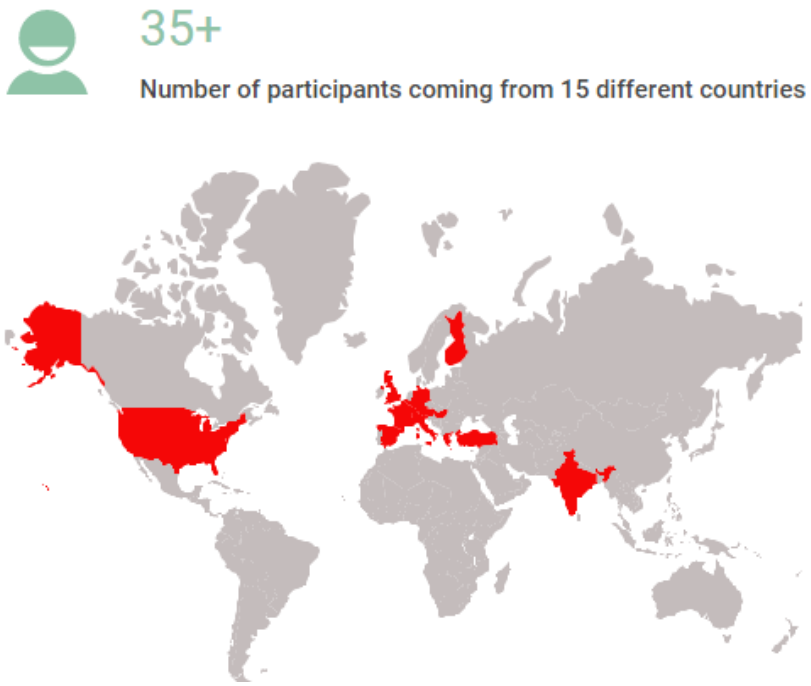

Figure 1: Attendee profile

There was a great deal of interest in the workshop: we received 44 registrations (the complete list can be found at https://goo.gl/V3PXqk), and more than 35 final participants, of which about 55% were from respective mPlane partners and about 45% from Universities, Research Centers or companies in Europe.

Below, some insights into the participants backgrounds:

- 65% of participants are within the 34-41 age range.
- 53% of participants attend more than two conferences per year.
- 81% of participants have Master's or Doctoral degrees.
- 40% of participants are Researchers or Professors.
- 40% of participants are Network Planning or Service Quality Engineers.

3.2.2 Agenda of the event

The following presentations/demos were part of the final demonstration workshop organized in Heidelberg:

- **Overview of the mPlane project**
  M. Mellia, mPlane General Project Coordinator
- **mPlane architecture and principles**
  B. Trammell, WP1 Leader (Use Cases, Requirements and Architecture)
- **Keynote: "Is measurement still an afterthought?"**
  Dr. Balachander Krishnamurthy, AT&T Labs – Research
- **mPlane live demo**
  F. Invernizzi, WP5 Leader (Integration, Deployment, Data Collection, Evaluation)
- **Keynote: "Content distribution on next generation cellular networks"**
  Prof. Fabián E. Bustamante, Northwestern University
- **mPlane components: probes**
  Á. Bakay, WP2 Leader (Programmable Probes)
- **mPlane components: repositories**
  M. Milanesio, WP3 Leader (Large-scale data analysis)
- **mPlane components: reasoners**
  P. Casas, WP4 Leader (mPlane Supervisor: Iterative and Adaptive Analysis)
- **A framework for the historical analysis and real-time monitoring of BGP data**
  Dr. Alberto Dainotti, Center for Applied Internet Data Analysis (CAIDA)
- **Demonstration plans and architecture**
  A. E. Kahveci, WP6 (Demonstration)
- **Use case demonstrations**
  Use case leaders

3.2.3 mPlane live demo

The purpose of this demo was to show a live, running example of an mPlane architecture application, on top of the integration system built in the context of WP5 activities.

The demo architecture ran on little hardware appliances physically present and shown at the demo stand, with probes (TI DATI, Tstat, pinger, HTTP latency), a supervisor and a reasoner, representing a complete and plausible setup (Figure 2). On startup, the WP5 reasoner (Figure 3) builds a static

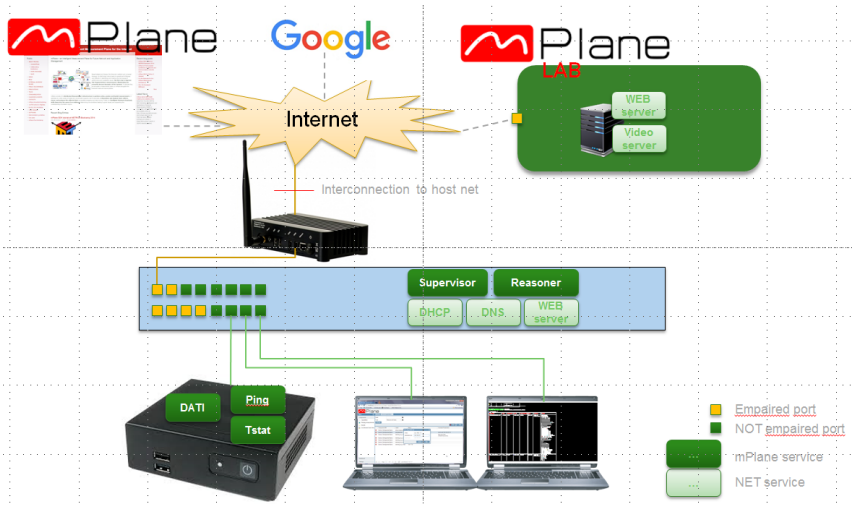

Figure 2: WP5 demo architecture

network topology graph based on a JSON description, calculates shortest path trees from any point to any other point in the network, communicates with the supervisor using the mPlane protocol in order to discover available measurements and finally maps measurements on the network graph.

After this startup phase, the reasoner iteratively issues measurements to probes by means of mPlane specifications in order to keep samples in a circular array per LAN (graph nodes) and periodically computes statistics of the network status.
The result is a JSON description of the network status that can be visualized in the supervisor GUI.

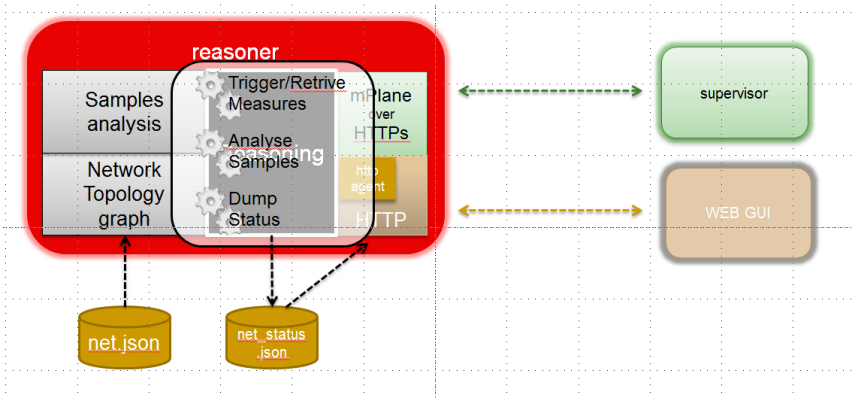

Figure 3: WP5 demo reasoner

3.2.4 Use case demonstrations

During the workshop, partners demonstrated the capabilities of the following use cases:

- Estimating content and service popularity for network optimization
- Passive content curation
- Mobile network performance issue root cause analysis
- Quality of Experience for Web browsing
- Active measurements for multimedia content delivery
- Anomaly detection and root cause analysis in large-scale networks
- Verification and certification of service level agreements

Following is a description of what has been demonstrated for each use case and the received input from the participants.

Estimating content and service popularity for network optimization

This demo presented a method to accurately estimate the popularity of online contents; specifically we looked at estimating the popularity of web-pages and online videos. The demo showed how these estimations can be used to optimize the network by reducing the network load, and enhance the user experience through selecting which contents to cache early.

Briefly, the content and service popularity modules built on the data collection tools provided through the Tstat interface. These provide the data which is used to build models (offline phase), and to later on evaluate the popularity of contents for selective caching (online phase). The data for the offline phase is provided via a database interface, while the online data is received via the Tstat live traffic capture interfaces.

In the demo, we showed how the models built using previously captured web content, and web video traffic can be used to optimize the performance of networks. Specifically, the selective caching methods in the demo enable network operators to reduce the load on their networks, and the latency of user traffic by selectively caching content based on its expected future popularity. Figure 4 shows the user interface for the content popularity estimation and the service optimization outputs.

Feedback during the event was positive. Observers showed interest in the underlying algorithms, and the use case of deploying such caches at the very edge of network where resources are limited, received support and interest. The people visiting the booth were positively surprised by the accuracy of the prediction provided by the demo. Some of them also raised questions about scalability issues and computational complexity of the algorithms. One of the visitors was concerned about the possibility of having a prediction model built on YouTube traffic data to work on generic HTTP-delivered content too (such as the content extracted by WeBrowse). We explained despite their different nature, often contents exhibit very similar popularity patterns. We also clarified that building specific models tailored towards specific content is not complicated.

Figure 4: Content popularity estimation and cache performance reporting.

Passive content curation

WeBrowse was demonstrated during the mPlane final workshop. The demo showed how the mPlane framework can be used to instantiate and launch the content curation use case.

The output of the demo is two fold. It first shows the WeBrowse web page which features a col-

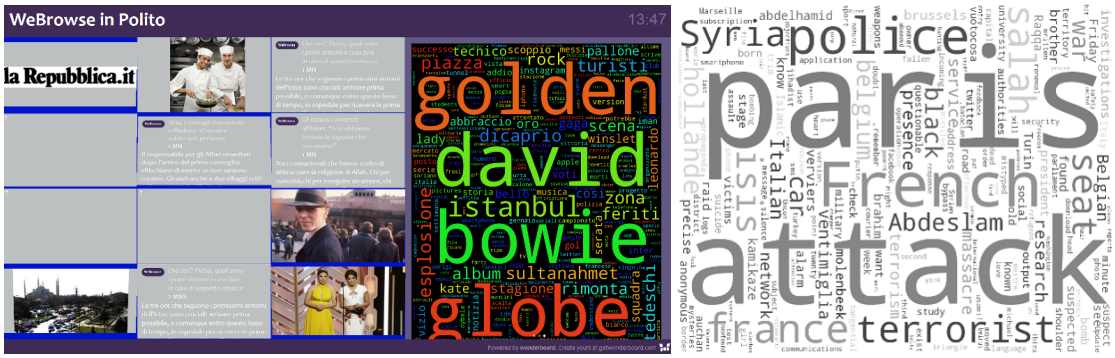

Figure 5: WeBrowse board and wordcloud

lection of the hottest content on the web at the moment. The second shows how the reasoner can provide insights about most popular web content in different time periods and presents them in terms of both most popular web pages as well as most popular topics discussed on the web in the form of a word cloud. Figure 5 a wordcloud of the topics featured on the web the day following the Paris attacks (November 2015).

People attending the booth were all positively surprised by the amount of information and content we could promote by inspecting traffic flowing in the network. Some visitors, not familiar with

WeBrowse and the methodologies it implements, were concerned about the privacy issues. We explained a little how WeBrowse solves these issues and this triggered an interesting conversation about the future of WeBrowse in fully-encrypted traffic scenarios.

Mobile network performance issue root cause analysis

In this demo, we demonstrated the ability of the mPlane architecture to facilitate the root cause analysis (RCA) of problems of video delivery. For this purpose, a set of mobile devices (Android phones and tablets), Wi-Fi access points and video servers were set up. In the demo, the participants were free to watch a video while we injected a set of impairments. While the video was playing the passive probes that were developed for the project collected a set of measurements from all network layers and uploaded the data to the repository.

An example is shown in Figure 6. Finally, the reasoner automatically detects the video session and

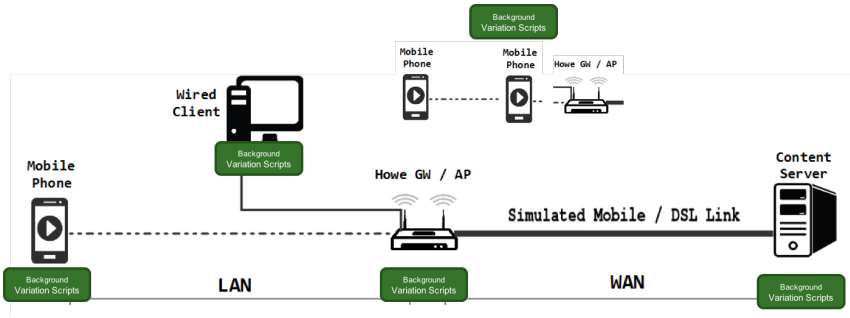

Figure 6: Mobile RCA demo architecture.

attempts to estimate the quality of experience and the root cause of a problem. The data can be requested using the integrated supervisor and the reference implementation. The last step is to visualize the result of the machine learning estimation as shown in Figure 7.

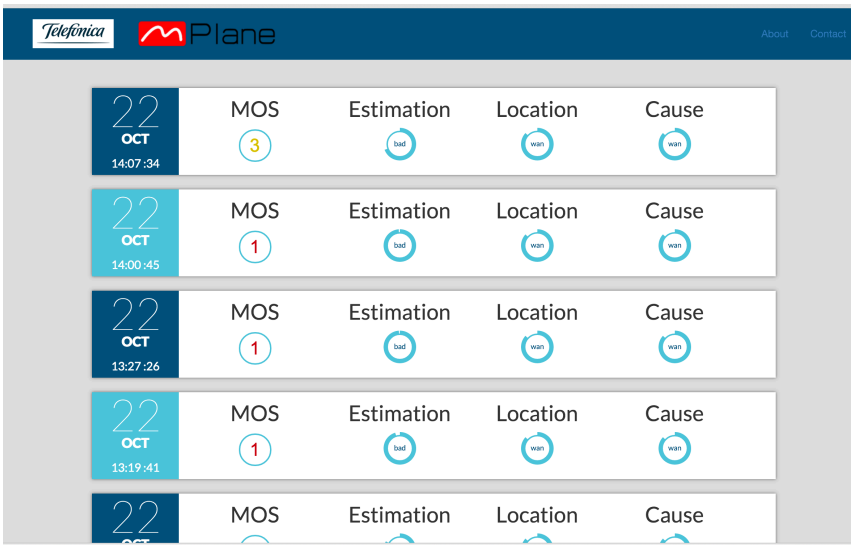

Figure 7: Mobile RCA demo visualization.

Overall, people were positive about the accuracy and the speed at which the classifier was able to predict the video quality. People showed particular interest in the features used, how do we measure them and what kind of machine learning algorithm has been used.

During the review meeting, the feedback for the mPlane mobile probe was positive mostly due to the fact that we used the iterative approach and the multiple vantage points that the mPlane architecture provides. More specifically, participants were impressed with the possibility of allowing different entities to manage different vantage points and that the architecture allows to have distributed troubleshooting across different supervisors. Also, during the demo, people were interacting with the mobile devices and trying to create video stalls and see if these were detected by the machine learning component. Finally, participants were interested in what was running at each hardware component (router, video server, access point, mobile phone).

Negative feedback had to do with the fact that not much information was given of the visualisation related to what network features where used to detect the problem (something that was not really our target as the GUI was mostly addressed to end-users and not to network experts).

Quality of Experience for Web browsing

The use case was demonstrated showing two browsing sessions carried by the Firelog probe.

In the first web browsing session, no impairments were added and we showed the overall process taking place during the measurement. In particular, we explained the active measurement phase, in which the Firelog probe performs active measurements, namely *ping* and *traceroute* against the IP addresses collected by the browser.

In the second experiments (with impairments), we showed that even if the page load time was only slightly touched, the diagnosis algorithm was able to identify a problem in the local area network (in which the impairments were inserted), demonstrating the importance of active measurements in order to have the complete picture of what is really happening *under the hood* in a web browsing session.

The integration within the mPlane architecture was part of the discussion, and we underlined that, with the Firelog probe fully integrated, it would be straightforward to integrate other measurements from different probes to further enhance the diagnosis process.

Active measurements for multimedia content delivery

This use case demonstrated a consistent and cooperating set of mPlane components and their communication fully in compliance with the mPlane protocol.

The problem domain here is monitoring and root cause analyzing of video content delivery networks that work "over the top" of the TCP/IP protocol set, i.e. consists of content servers delivering media streams over the classic unicast routing mechanisms. The latest of these protocols (HLS, MPEG-DASH, etc.) all support some level of adaptivity, i.e. the server offers content in various qualities (and bandwidths), and if a client experiences poor download bandwidth for a certain quality it can switch back to a lower level which downloads faster.

The demonstration includes a redundant pair of content servers in the FASTWEB data center and several customer locations (i.e. emulated FASTWEB residential subscribers). There are various probes deployed either in the DC, in the provider network, but also in the subscriber premises of

selected "customers". The demonstrated scenario starts with some constant basic-level monitoring of probes, i.e. occasional downloads of varying content by the residential probes, to verify availability and sufficiency of bandwidth toward any of the content servers.

All these "routine" verification data is collected by the repository (EZ-Repo), which stores recent measurements for about 24 hours. The repository is periodically (every 60 seconds) queried by the Reasoner, for preset, topology-aware evaluations of service quality. Whenever some abnormal or degraded condition is indicated by any of these queries, the Reasoner automatically configures "diagnosis tests" on further probes. These measurements are more focused and more frequent than the routine tests. Once the result of these measurements is also available, the Reasoner will decide which of the hypotheses have been supported and presents the results on its "service diagnosis" dashboard.

The use case demo makes extensive use of the supervisor GUI as shown in Fig. 8, which is a complete, web-based interface for presenting the state and history of mPlane components' operations: one can view capabilities, pending and finished specifications, and also the results of those specs. In addition, it is also possible to present single or multiple results in chart views, also with live updates of new data. The mPlane GUI also incorporates a "Dashboard" capability where users can configure customized screens that present system status and history using well-known visual tools like charts, tables, gauges, maps and other indicators. As for the probe side, the multimedia con-

Figure 8: Multimedia content delivery UC dashboard

tent delivery UC primarily relies on the OTT Probe developed in WP2. It is an mPlane probe which receives OTT content URLs as part of the specifications, and starts the downloads from the URLs specified. While doing this, the probe emulates an adaptive client, i.e. switches between content qualities as required.

Visitors were generally interested and impressed with the demonstration, which essentially relied on the mPlane Supervisor GUI, and also in the probe's local display. Several people wanted to know, how the probes assess the service quality of streams received, and what are the key mertics defined.

On the other hand, the more generic, more sophisticated and more exciting mPlane components, i.e. our RC1 reasoner and the EZ repository did not get comparable attention. One visitor noted that he does not yet see these components scalable and generally applicable outside of the demonstrated monitoring task. While scalability was indeed not addressed yet, we see the modules' general applicability quite straightforward for a wide range of other diagnostic problems as well.

After the final demo (but before the Year3 review in Milano) we added a few more dashboard panels to highlight this and make the reasoner and repo components more visible. Most prominently, a new panel to demonstrate the reasoning "waterfall" timeline was introduced. We found that this improved demonstration was suitable for explaining the concepts of the UC and the role and interactions of the components involved.

Anomaly detection and root cause analysis in large-scale networks

In this use case demo, we have shown two different application scenarios.

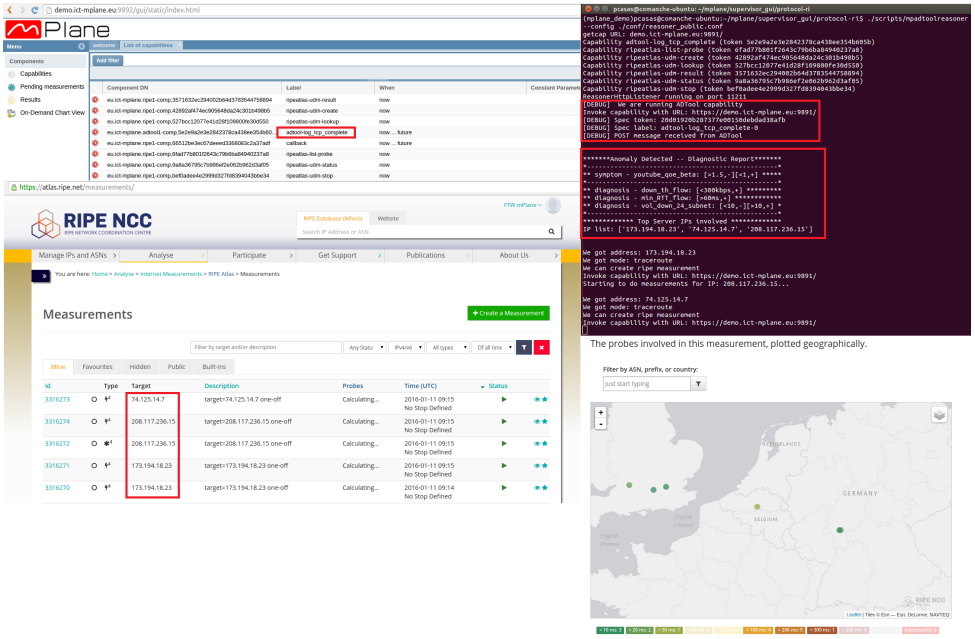

Figure 9: Detection and diagnosis of anomalies in YouTube - demo.

In the first part of the demo, we have made a proof of concept for the mPlane Anomaly Detection modules, showing how these can effectively detect anomalous behaviors related to QoE-based performance metrics, and help in the root cause analysis investigation. This part of the demo was based on historical YouTube traces.

The second part of the demo showed how to use the mPlane mpAD Reasoner to orchestrate the live collection and analysis of passive and periodic active measurements, and to trigger on the fly further active measurements. In particular, this part of the demo used active measurements returned by DisNETPerf (periodic/continuous, based on RIPE Atlas, see [6]), and the RIPE Atlas mPlane proxy instantiation (reactive, on-demand). The presentation and introduction to the different parts of the demonstration was done through the aid of posters, particularly describing the functioning of the anomaly detection modules and the reverse traceroute capabilities provided by the DisNETPerf module.

Figure 9 shows a snapshot of the different steps followed during the demo (note: for practical and reporting reasons, the dates shown in the figures do not correspond to the date of the demo, but the results are exactly the same), including part of the results in the different GUIs. Starting on the upper left corner, the mPlane public Supervisor running at FASTWEB premises in Mi-

lano (http://demo.ict-mplane.eu:9892/) shows the main registered capabilities which compose this demo (anomaly detection modules and RIPE Atlas integration modules). The upper right corner depicts the instantiation of the Anomaly Detection Reasoner, showing in particular the execution of the Anomaly Detection capabilities, the detection of an anomaly impacting YouTube's QoE, a list of IP addresses of some of the involved YouTube servers, and the subsequent instantiation of active measurements on these IPs using DisNETPerf and the RIPE Atlas framework. The lower part of the figure shows some of the instantiated active measurements on the aforementioned list of YouTube server IPs, using the RIPE Atlas measurement GUI. In particular, the lower left corner shows the instantiation of traceroute measurements towards these IPs, as well as the starting phase of the DisNETPerf module [6], in which the YouTube IP address 208.117.236.15 is pinged from multiple distributed RIPE Atlas probes (selected based on topological-based notions). The lower right corner of the figure actually shows these RIPE Atlas probes on a map, including a color code reflecting the RTT values obtained from the ping measurements towards the selected YouTube IP.

The feedback received during the demonstration was very positive and encouraging. Attendees recognized the flexibility and virtues of the mPlane approach: in particular, they recognized how easy it was to run this use case with so many geo-distributed components, including the mPlane public Supervisor running at FASTWEB premises in Milano, commanded by the AD Reasoner running locally in Heidelberg on an end-user laptop, detecting anomalies on the traffic monitored at PoliTO premises in Turin and instantiating new measurements on the fly from geo-distributed probes in Frankfurt, London, Paris, and so on using the RIPE Atlas distributed framework. All this is done without having to access a single piece of private ISP data, which was also highly appreciated and recognized by the attendees during discussions.

Verification and certification of service level agreements

Four different partners (FUB, NEC, FASTWEB, Telecom Italia) were involved in the demonstration of the use case and the results confirmed the reliability of the mSLAcert approach for the SLA verification and certification. In particular, such a test verified that the mSLAcert probe (based on UDP test) allows us to correctly measure the line capacity of links also having a " high bandwidth delay product", that are the typical conditions where TCP tests are not reliable.
In Fig. 10, we present the demo network including all four operating nodes. Use case components (SLA Server and SLA Agent) were distributed to the following premises: FUB, NEC, FASTWEB and Telecom Italia.

All the components of the use case were integrated with the public Supervisor, located at FASTWEB premises where network impairments were also applied. For sake of brevity, we will report only two tests that were executed. First, we describe the SLA test based on TCP.

As it can be seen in Figure 11, we measure the maximum capacity as 50 Mbps, while the line capacity that is reserved on both sides is 100 Mbps. Therefore, the SLA test based on TCP shown in Figure 11 does not allow us to measure the maximum available line capacity, under the introduced impairments (a delay of 75ms per direction). This is the typical case where it is required to adopt a UDP test in order to execute SLA measurements. Therefore, the complete tests of an SLA, with TCP and UDP are shown in Figure 12.
As can be seen in Figure 12, that the UDP test measures a maximum capacity of 92 Mbps, while the TCP test only measures 43 Mbps. As can be seen in Figure 12, the connections show a relevant

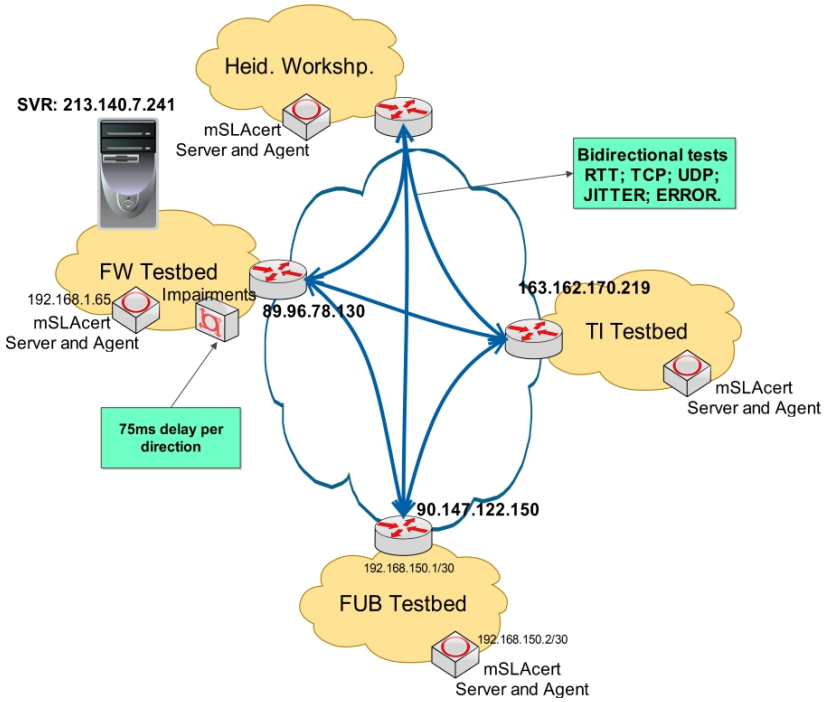

Figure 10: SLA demo architecture.

jitter and this impairment induces other limits in the reliability of the measurements based on TCP tests. Conversely, UDP tests are independent of RTT and jitter and therefore UDP can be assured as the best approach for the line capacity evaluation.

The results shown in this test confirmed the importance of the mSLAcert probe in the evaluation of the line capacity related to the access network (from the central office to the user modem), also in the presence of high RTT and jitter (Fig 12).

During the demo, we received positive feedback. In particular, the participants were pleased to see the graphical interface which presented the difference among the UDP (constant at 99Mbps) and TCP (lower than 50 Mbps) throughput measurements.
The critics that we received were mostly about the limitation of the tests. The participants were eager to run measurements also for different line capacities and technologies like 4G; those tests were experimented only after the end of the project.

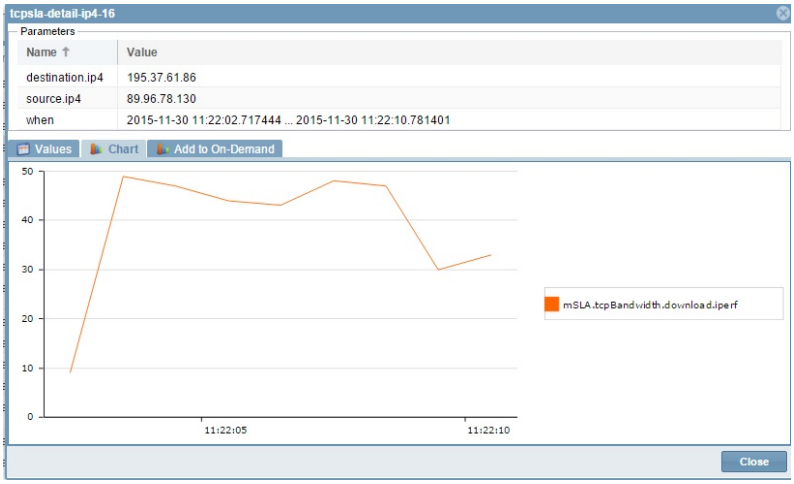

Figure 11: SLA test based on TCP, RTT of 150ms

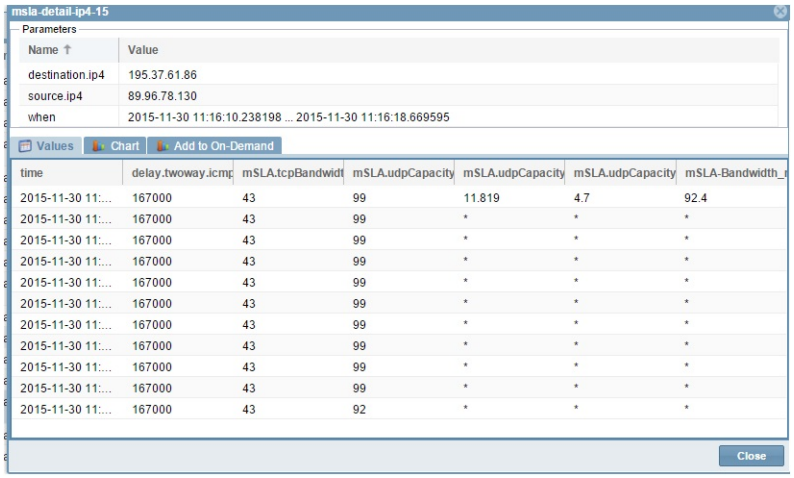

Figure 12: SLA test with TCP and UDP.

3.2.5 Other activities demonstrated

During the workshop, partners also demonstrated other mPlane activities. What follows is a description of these other activities demonstrated in Heidelberg and the received input from the participants.

Anycaster

Use of anycast IP addresses has increased in the last few years: once relegated to DNS root and top-level domain servers, anycast is now commonly used to assist distribution of general purpose content by CDN providers. Yet, most anycast discovery methodologies rely so far on DNS, which limits their usefulness to this particular service. This raises the need for protocol agnostic methodologies that should additionally be as lightweight as possible in order to scale up anycast service discovery. Our anycast discovery method allows for exhaustive and accurate enumeration and city-level geolocation of anycast replicas, with the constraints of only leveraging a handful of latency measurements from a set of known probes. The method exploits an iterative workflow to enumerate (optimization problem) and geolocate (classification problem) anycast instances. The method is so lightweight and protocol agnostic that we were able to perform several censuses of the whole IPv4 Internet (during March 2015). A web interface (http://perso.telecom-paristech.fr/~drossi/anycast) with our census results was shown in the final workshop in Heidelberg, see Figure 13.

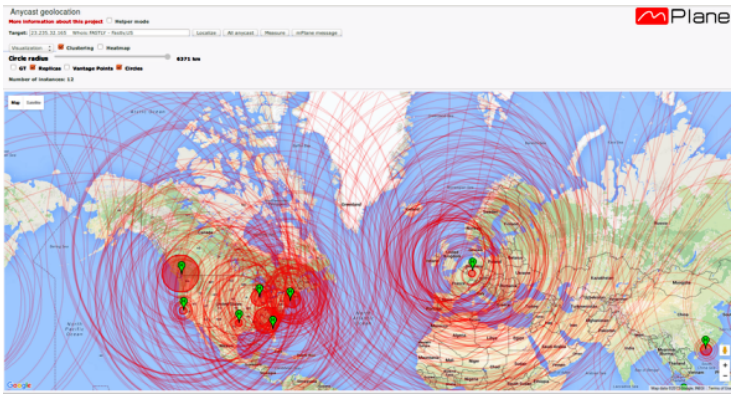

Figure 13: Anycaster demo snapshot

GLIMPSE

To show that the mPlane architecture can be used outside of one of the use cases specified by the consortium, GLIMPSE was demonstrated at the workshop in Heidelberg. GLIMPSE is an active network measurement tool intended to be installed on an end-user devices at home or on a mobile device such as a smartphone or tablet. It can execute centrally-defined measurements delivered by a supervisor as well as measurements the users schedule themselves to be executed on their devices. It comes with a web-based dashboard, where users can visualize the executed measurements and define their own measurement schedules. All of this was shown as part of the demo, including graphs of ongoing measurement campaigns. Part of the demo were also live probes measuring over a 3G home gateway.

The demo walkthrough started with an overview of where GLIMPSE can be deployed together with

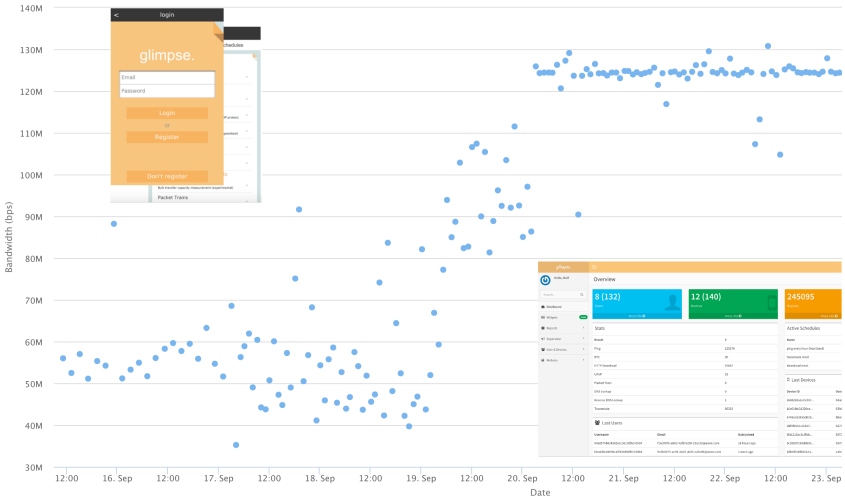

Figure 14: GLIMPSE demo snapshot

some of the data gathered so far during the beta-testing phase of the software (see Fig 14) The feedback received was very positive, the graphical results shown lead to interesting discussions. As the main focus of GLIMPSE is the end-user, some really helpful insights and suggestions from users but also from employees of ISPs could be gathered to further improve GLIMPSE in future releases.
This demo was also the basis for the previously described demonstration in Berlin as part of the 9th EAI International Conference on Performance Evaluation Methodologies and Tools.

ECN path transparency

The ECN Path Transparency demonstration, built on ETH's Pathspider tool, demonstrates the benefits of direct integration of tools over the mPlane protocol for research purposes. Pathspider is built as a hybrid component-client with its own web-based user interface. Its purpose is to test connectivity to an arbitrary remote target on the Internet, both with and without a protocol feature enabled, in order to determine whether connectivity is impaired for that feature. In order to separate endpoint- from path-related connectivity impairments, Pathspider uses mPlane to coordinate tests from multiple vantage points simultaneously: if the feature causes connectivity impairment from all vantage points, the impairment is presumed to be at or close to the target; if only from some vantage points, the impairment is presumed to be on-path.

For purposes of the demo, TCP Explicit Congestion Notification (ECN) was tested, with a client/webui in Heidelberg, and components on five different VMs running at DigitalOcean in Amsterdam, London, Singapore, San Francisco, and New York.
The demonstration used Pathspider's web user interface to illustrate the multiple stages of this measurement. First, the user selects a set of targets (on the order of hundreds), either web servers (which will be resolved to IP addresses at a single vantage point, via a DNS resolution capability) or peer-to-peer nodes (based on a BitTorrent DHT crawler capability). The target IP addresses are then sent back to the client, which distributes them to the components. Each component tests connectivity (via a connectivity status capability) to each of the targets, both with and without ECN, determining any ECN-dependency on connectivity. The results are then sent back to the client, which

displays aggregate statistics (how many targets are ECN-safe, how many have endpoint-linked ECN connectivity dependency, how many have path-linked dependency, and so on) as well as allowing the user to drill down into routes to a specific target.

If this option is selected, the client instructs each component to find packet modifications along the path to the target (via the scamper/tracebox component built by ULG, used unmodified in this demo), displaying the traceroute from five sources to one target, and allowing the user to inspect each node (via an external hyperlink to the RIPEstat facility).

3.2.6 Comments, general impressions and feedback from the audience

This section describes the overall impression of the participants, as well as how the mPlane project and its results were received. We created an online questionnaire in order to obtain information from participants on how they perceive the mPlane project and its outcomes.

After concluding the mPlane workshop, we began collecting feedback from the participants via the online survey, with the aim of capturing the opinions of those who had attended, and see whether the algorithms and the use cases presented during the workshop were found to be useful or not. In addition to the many positive informal responses received during and after the workshop, we received 21 replies to the survey, with an average of about 50% survey completion.

We will begin by highlighting some of the more significant results, followed by a detailed per-question breakdown of the survey.

Feedback: mPlane workshop Heidelberg Demonstration Results (based on the online survey responses received):

- 62% thought that the overall presented solution is very useful
- 80% thought that they would very likely or somewhat likely use the overall presented solution
- 76% thought that the use cases presented are positive or very positive

Attendees:

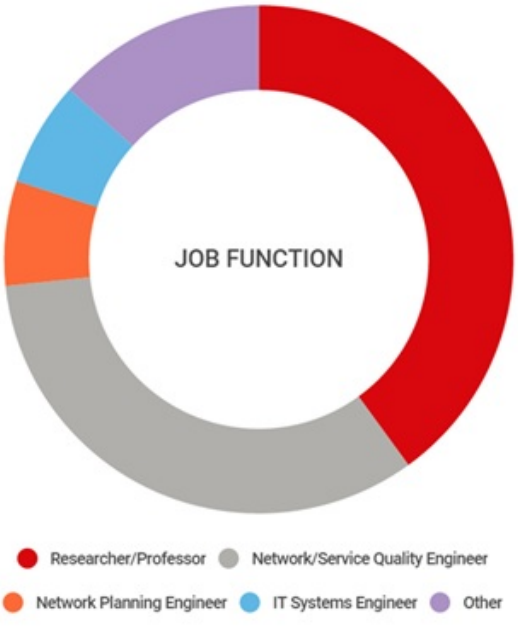

Percentage of survey participation:

~50%

Percentage of respondents who attend at least two conferences a year:

53%

Overall project feedback:

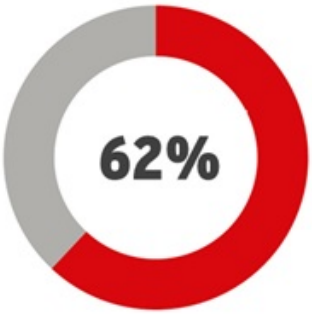

Percentage of respondents who find the overall presented solution very useful

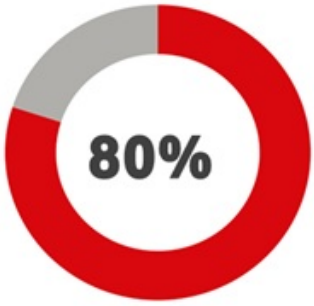

Percentage of respondents who think they would likely use the overall presented solution

Use case feedback:

As seen from the responses, the overall feedback towards the project's results was very positive and encouraging.

In addition to the positive feedback, it is also important to acknowledge some criticisms from the audience.

As this workshop was the only one at which all of the use cases were demonstrated (ten in total), the agenda was densely packed, and some of the participants found it difficult to understand and experience all of the use cases because of sheer volume.

Moreover, as stated earlier, the workshop hosted people with a variety of backgrounds and levels of experience. While everyone, regardless of their experience, was appreciative of the flexibility of the demonstrator, as well as the graphical interface, some end users, mainly students with little to no experience in the field, found it difficult to set up the configuration of the use cases and the platform integration. It should be noted that, we have a dedicated section [1] on the mPlane official website that explains step-by-step installation/execution of use cases shown during the demonstration events so that end users could run their preferred use cases.

mPlane Heidelberg Demonstration Feedback Reports
The details of the feedback received for each question can be found below:

Question 1

How useful did you find the overall presented solution?

Question1

Options	Quantity	%
Very useful	13	61.90%
Somewhat useful	8	38.10%
Not very useful	0	0.00%
Useless	0	0.00%
Total	21	100%

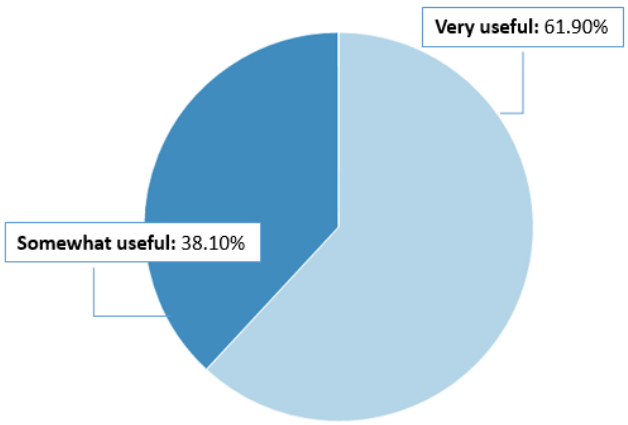

Question 2

How likely is that you would use some of the presented solution?

Question2

Options	Quantity	%
Very likely	10	50.00%
Somewhat likely	6	30.00%
Neutral	4	20.00%
Somewhat unlikely	0	0.00%
Very unlikely	0	0.00%
Total	20	100%

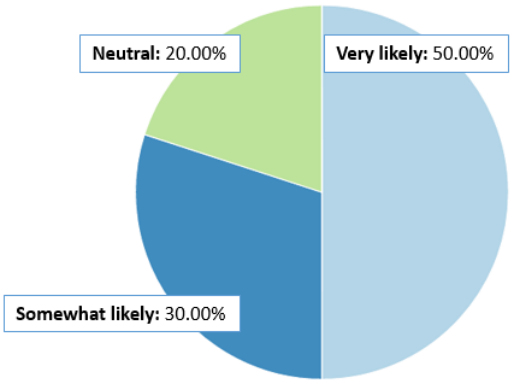

Question 3

How would you rate the following use cases

Question3

Matrix Chart

	1	2	3	4	5
Estimating content and service popularity for network optimization	0	0	4	7	5
Passive content curation	0	2	2	8	4
Active measurements for multimedia content delivery	1	0	2	7	6
Quality of Experience for Web browsing	0	0	5	7	4
Mobile network performance issue cause analysis	0	0	4	7	5
Anomaly detection and root cause analysis in large-scale networks	0	0	3	6	7
Verification and certification of service level agreements	0	1	3	7	5

Question 4

What is the highest level of education you have completed?

Question4

Options	Quantity	%
Doctoral degree	6	37.50%
Master's degree	7	43.75%
Bachelor's degree	3	18.75%
Other	0	0.00%
Total	16	100%

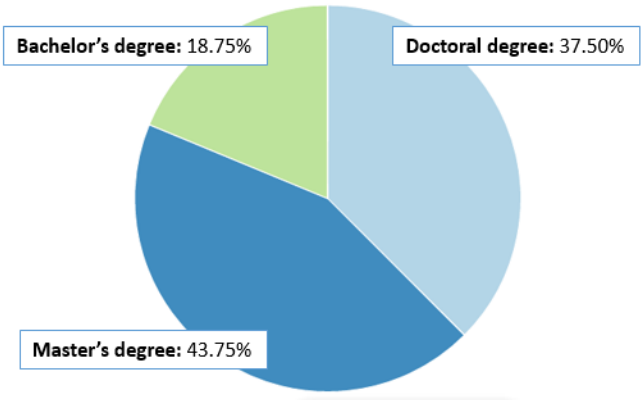

Question 5

What is your expertise about the topic of the mPlane project?

Question5

Options	Quantity	%
Network Administrator/Manager	0	0.00%
Network/Service Quality Engineer	5	33.33%
Network Planning Engineer	1	6.67%
IT Systems Engineer	1	6.67%
Researcher/Professor	6	40.00%
Other	2	13.33%
Total	15	100%

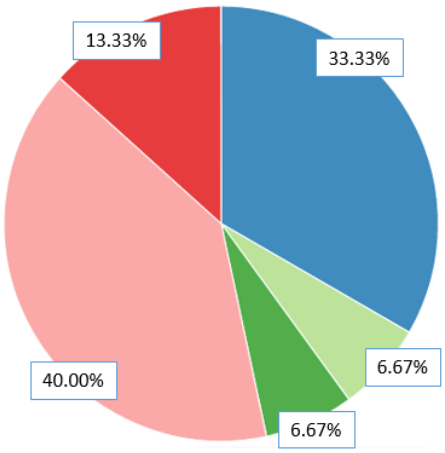

Question 6

How many conferences per year do you attend on average?

Question6

Options	Quantity	%
1	7	46.67%
2-4	5	33.33%
5+	3	20.00%
Total	15	100%

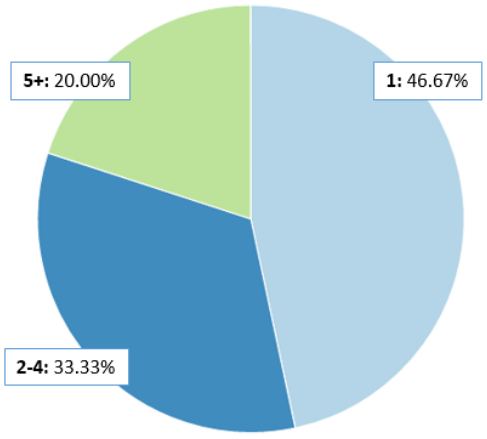

Question 7

Select your Age

Question7

Options	Quantity	%
18 - 25	1	6.67%
26 - 33	4	26.67%
34 - 41	9	60.00%
42 - 55	1	6.67%
55+	0	0.00%
Total	15	100%

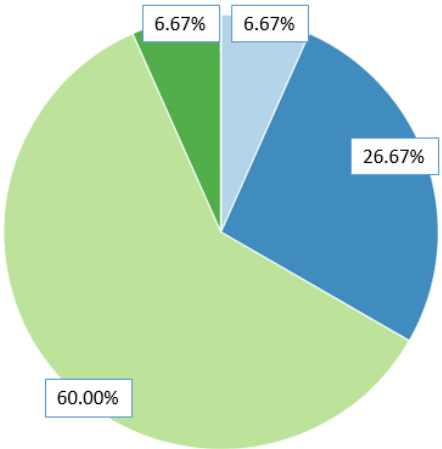

Below follows a set of photos from the mPlane workshop and live demonstrations to get an impression of the projects demonstration activities.

Figure 15: A presentation and the audience at the workshop

Figure 16: Attendees at the live demo

Figure 17: Use case live demos

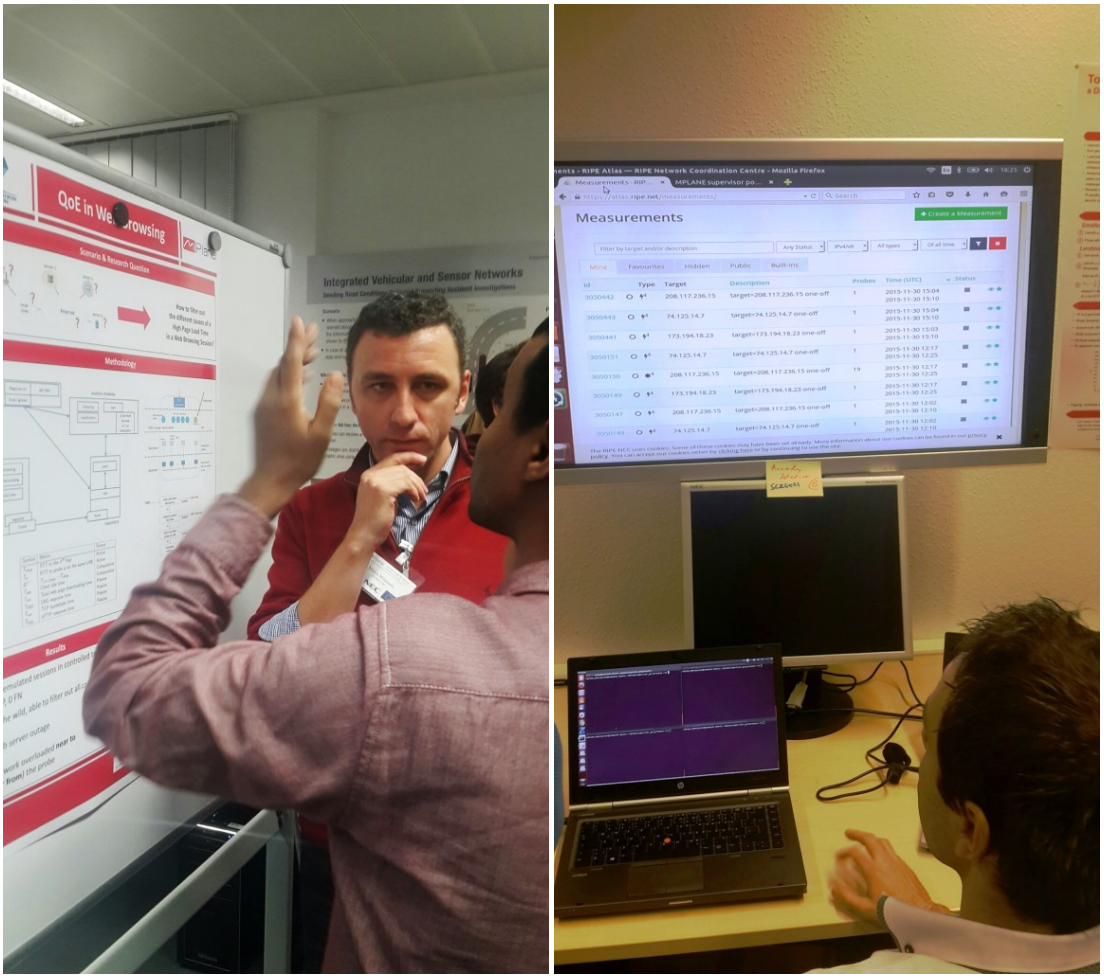

Figure 18: Use case live demos

4 Conclusions

The mPlane partners have organized and participated in various events, where they have introduced the mPlane project and demonstrated the capabilities of the built system.

During the demonstration events, we received very positive reviews and feedback from the respective participants, and mPlane has put great effort into properly disseminating the results of the project, namely through the mPlane workshop which took place on November 30, 2015, in conjunction with CoNEXT 2015.

In conclusion, the dissemination of the mPlane through demonstrating results life on the built software has been successful and very well received.

References

[1] https://www.ict-mplane.eu/public/demonstration-guidelines.

[2] JCSA15. https://www.afnic.fr/fr/l-afnic-en-bref/actualites/actualites-generales/9333/show/jcsa15-retour-sur-la-journee-du-conseil-scientifique-de-l-afnic-2015.html.

[3] D. Karrenberg. Le réseau mondial des sondes “ripe atlas”, au service de la communauté internet scientifique.

[4] D. Rossi. The mplane distributed measurement infrastructure: overview, insights & hindsight.

[5] Twitter. https://twitter.com/search?q=%23JCSA15&src=typd.

[6] S. Wassermann, P. Casas, and B. Donnet. Towards disnetperf: a distributed internet paths performance analyzer. In *The 11th International Conference on emerging Networking EXperiments and Technologies - CoNEXT 2015*, 11/2015 2015.