MONROE

Measuring Mobile Broadband Networks in Europe

H2020-ICT-11-2014
Project number: 644399

Deliverable D6.4

Second Year Report

Editor(s): Özgü Alay, Andra Lutu, Katarina Subakova
Contributor(s): Anna Brunstrom, Audun Fosselie Hansen, Giacomo Bernini, Vincenzo Mancuso, Min Xie

Work Package: 6 / Project Management
Revision: 1.0
Date: July 3, 2017
Deliverable type: R (Report)
Dissemination level: Confidential, only for members of the consortium (including the Commission Services)

European Commission
Horizon 2020
European Union funding
for Research & Innovation

Abstract

This report summarizes the progress of the MONROE project throughout its second year of activity.

Participant organisation name	**Short name**
SIMULA RESEARCH LABORATORY AS *(Coordinator)*	SRL
CELERWAY COMMUNICATION AS	Celerway
TELENOR ASA	Telenor
NEXTWORKS	NXW
FUNDACION IMDEA NETWORKS	IMDEA
KARLSTADS UNIVERSITET	KaU
POLITECNICO DI TORINO	POLITO

1 Activities Overview and Progress Update

In the DoA we defined the following main objectives for MONROE:

1. To build an open and large-scale measurement and experimental platform, targeting MBB networks and WiFi connectivity, distributed over multiple European countries, with multi-homing capabilities.

2. To operate this large-scale platform by providing both maintenance and external user support.

3. To use the platform for the identification of key MBB performance parameters, thus enabling accurate, realistic, persistent and meaningful monitoring and performance assessment of such networks.

4. To achieve a user-oriented closed-loop system design in which the experimental platform is open to external users, and where users are incorporated early on in the experimental design process.

5. To provide Experiments as a Service (EaaS), thus lowering the barrier for using the platform to external users, by providing well-documented tools and adjustable, flexible, high-level scripts to execute experiments, collect results, and analyze data.

6. To develop models for sustaining and extending the platform and its usage beyond the point when the awarded project funding ends.

Next, we describe the work we carried out during the second year of project activity towards the achievement of each of the above-mentioned objectives.

Objective 1: To build an open and large-scale, transnational experimental platform

During the second year of the project, we have focused on building and deploying MONROE nodes. Within this process, we worked towards refining and improving the system design in general and the MONROE node design and implementation in particular. We relied on and incorporated the feedback received from the first MONROE open call users. Additionally, we built on the testing experience of the project partners while deploying the MONROE nodes in Norway, Sweden, Spain and Italy. After a six-month period of operating the prototype platform, we started experiencing significant problems with the USB hub and the MiFis. Especially, the NAT settings and battery issues experienced with the MiFis forced us to re-consider the initial node design. After extensively searching for new hardware components to replace the MiFis and the USB hub, the consortium came up with multiple alternative node designs. After spending significant amount of time to extensively evaluate these alternatives (e.g. testing the new node design alternatives in the lab as well as on the trains and busses), the consortium agreed on the new node design where the USB hub and MiFis are removed and replaced by PCI express modems in a dual-APU setting. We provided the details of the problems we experienced with the initial node design, alternatives that are evaluated and the details of the new node design and implementation in Section 3.1.

The MONROE system design maintains the same general structure we defined in the first year of the project. We have configured the upgraded MONROE node to reuse as much as possible the software components we developed throughout the first year of project activity. Moreover, we carried out the hardware upgrade such that it is transparent to experimenters that are not interested in hardware details. In this way, users can proceed with their experiments as if the platform components were unchanged. Though the new node design translates into more robust hardware and overall increased platform stability, the upgrade of the nodes' implementation to use PCI express modems affected the total number of nodes we are able to build (from 250 nodes planned, to 150 nodes with the new design). This is due to the new node with dual-APU integrating two main APU2 boards, each able to accommodate two different PCI express modems, hence being more expensive. In light of this upgrade, we have refined the interaction of the user access and the scheduler systems, as well as management routines and the orchestration of data importing and database operations. The system upgrade is transparent to the external users, and each individual APU2 board can individually cater to different experiments (see Section 3.1 and Appendix II for details on our experience and the platform upgrade).

During the second year of the project, we have finalized the MONROE implementation with help of our external experimenters. The final system implementation of the MONROE platform consists of seven different components: (i) the user access portal, (ii) the scheduling system, (iii) the management and maintenance system, (iv) the software running on the MONROE node (core components and experiments), (v) the repositories and data importer, (vi) the central database, and (vii) the online visualization tool. We documented these components in D1.3. Moreover, we open-sourced each component to the community[1]. Because of the hardware upgrade, the node deployment is delayed. We reported initial node design in D2.1 and the node deployment in D2.2. We will update these document upon finalization of the deployment of the upgraded nodes in June 2017.

With the submission of the deliverable on final system implementation (D1.3), the software part of the objective is achieved. Upon re-submission of deliverables on node design (D2.1) and node deployment (D2.2), this objective will be achieved.

Objective 2: To operate a transnational and large-scale platform

The operations of such a large-scale platform such as MONROE require maintenance activities, as well as user support. During the second year, we provided maintenance for the nodes and for all the software components of the system, as well as the measurement data. The maintenance activities involve both hardware and software maintenance, including the replacement of the physical nodes when necessary. We defined and reported the maintenance routines in D4.1. We ran weekly maintenance meetings to follow them closely. We also developed and implemented monitoring tools and a robust node recovery method (D2.3) to minimize local intervention, which would be particularly expensive. Beside the testbed maintenance, we provided direct support for external experimenters and testbed users. For the first and second open call users, each user has been assigned a patron that is responsible for providing immediate support. We found this approach effective towards resolving any issues related to the use of the platform that may come up during the experiment. Furthermore, we activated a mailing list for the open call users in an effort to build the MONROE community, encourage them to interact, exchange experience and help each other. Part of the user support consisted in preparing a user manual that is easy to follow and up-to-date with most recent information about the MONROE testbed. We also collected feedback from the first open call users and took necessary actions to update the platform capabilities to meet the user's needs. We summarized the user manual and the initial user feedback in D5.1. The utilization of the platform throughout the second year was high, with more than 70,000 experiments scheduled to run on the platform. The numerous paper submissions from both the consortium and the external users is a good measure of both the utilization and utility of the platform (for a full list of publications see Section 8).

The project is on track with achieving this objective despite the delay due to the hardware related problems.

Objective 3: Identification of key performance parameters

One of the main objectives of MONROE is to measure and monitor the MBB networks accurately and fairly. To achieve this, it is crucial to identify the metrics that accurately capture the performance and the conditions under which these parameters should be measured. However, the parameters of interest might be different for different stakeholders. For example, regulators need connectivity, coverage and speed information collected from a third-party, independent platform to monitor how operators fulfill their obligations,

[1]MONROE on GitHub: https://github.com/MONROE-PROJECT.

and as a baseline for designing regulatory policies. On the other hand, operators are interested in operational instability and anomalies to identify problems in their networks. For end-users, QoS and QoE parameters are of paramount importance, while application developers need information about the underlying network to design robust services and protocols.

In the second year of the project, we focused on defining measurement methodologies and verifying them experimentally in order to accurately reflect the performance and reliability of MBB networks from the perspective of different stakeholders. We have discussed with external users and other stakeholders in MBB networks about which are the parameters they are interested in and how the measurements should be carried out. This resulted in a set of experiments that covers different use cases defined in D1.1, including coverage analysis, speedtest measurements, web and video performance analysis, middlebox detection and path support analysis. These experiments[2] are designed and implemented, and the results of the measurements analysis has already lead to multiple publications (accepted or under review). Along with these publications, we have provided sample dataset from these measurement results as Open Data through project's website as well as in public repositories such as the Zenodo[3] repository, to enable any interested party to analyze and extract value from these results.

Overall, the project is on track with achieving this objective.

Objective 4: User-oriented closed-loop system design

For a testbed to be useful to its stakeholders, the capabilities of the testbed have to meet a wide range of expectations. MONROE involved potential users from the early stages of system design. This approach allowed us to collect feedback and engage in open communication with the users we selected after the first MONROE open call for external experimenters. Their feedback was very useful in the design and the implementation process of the testbed. This was possible thanks to the fact that we have designed the MONROE nodes as small programmable computers and we have developed open software for managing the testbed, resulting in a fully programmable MONROE system. The operation of nodes is abstracted thanks to (*i*) the adoption of virtualization techniques such as the one offered by Docker containers, and (*ii*) the creation of a Fed4FIRE-compliant experimenter interface to the nodes, using open source resources to enable a user access portal.

Within the second year of project, we completed the selection process for external experimenters after the MONROE second open call. In total, we now have 27 external users, 12 selected in the first open call and 15 in the second one. First open call users had access to the prototype system and provided us feedback in order to improve the infrastructure. In the second year, the experience and feedback from the first phase external users (D5.1) have been used in upgrading and modifying the system design and implementation of the platform. Specifically, the impact of internal and external feedback on the system design is documented (D1.3) and is measurable in terms of design modifications between the prototype and the final release of the platform.

This objective is achieved with the submission of deliverables on final SW implementation (D1.3) and on the user manual (D5.1).

Objective 5: Lowering the barrier for external users

One of MONROE's main targets is to make the experimental platform and designed experiments open to external users from the community. However, learning and getting used to a new tool is always a time-consuming task and designing experiments, especially for large-scale complex systems such as MBB networks, requires experience and time. Therefore, one fundamental objective of MONROE consists in pro-

[2] https://github.com/MONROE-PROJECT/Experiments
[3] https://zenodo.org/communities/h2020_monroe/

viding comprehensive documentation for external users about how to use the platform. The user manual[4] has been prepared during the second year and continuously updated on a live repository made available to MONROE experimenters. Further, we have publicly provided in our public github repository several EaaS examples[5], including a template embedding high-level scripts to run simple experiments or fully developed experimental containers to run different measurements. In addition, some experimenters of the first open call have already started to release their experiment scripts and containers. The MONROE team takes care of integrating such contributions within a shared repository of EaaS examples and provides support to keep the examples aligned with the evolving platform.

In our efforts to communicate and cater for the needs of our external users, we have also arranged two user workshops for external partners. The first one took place in early June 2016, and the second one in October 2016. Throughout the second year, with the external users accessing the platform, we have been providing support for External Users both in terms of using the platform as well as designing their experiments.

This objective is achieved with the submission of deliverables on experimental software as EaaS (D3.1) and on the user manual (D5.1). The consortium is committed to continue sharing new experiments as EaaS for the external users and keeping the user manual uptodate for the remainder of the project.

Objective 6: Sustainability of the platform

Despite challenges with the hardware implementation that translated to delays in our initial timeline, partners are very optimistic and see the great value in using and supporting the MONROE platform. We have ensured that all partners (both consortium and external partners we added through the MONROE open calls) are dedicated to sustain and extend the platform. The sheer number of applications for the open calls (40 applications for the open call 1 and 50 applications for open call 2) and the feedback of selected partners in workshops prove the need of the community for an initiative like MONROE and the potential of the platform. MONROE open call project proposals have also highlighted various business models and sustainability plans for MONROE, resulting in seven viable ideas for business and sustainability (listed in Section 8.3).

Visibility in media, webpage and EC events have resulted in invitations to EU project proposals, interest from institutions and governmental bodies. Furthermore, the support for MAMI[6] components and the NEAT[7] library in MONROE nodes increase the interest of the community in MONROE. Following on up on the interest, we have analyzed competitors in order to identify key competitive advantages with MONROE. This will be input to a more focused work on business models during the first half of the third year. In Norway, MONROE concepts are already supported by two operators, a train company and a ministry. One goal will be to copy similar models in other European countries.

The project is on track with achieving this objective despite the delay due to the hardware related problems.

2 Work Package 1

Work Package 1 involved the design of the system's building blocks and the implementation and refinement of required components. During the second year, the effort has been dedicated to the Phase II of the implementation, evolving from a prototype to an operational transnational platform. Once the node's hardware selection was finished (MS 2) and Task 1.1 was completed at the end of the first year, the effort has been

[4] https://github.com/MONROE-PROJECT/UserManual
[5] https://github.com/MONROE-PROJECT/Experiments
[6] https://mami-project.eu/
[7] https://www.neat-project.org/

focused on following tasks:

- Task 1.2: Implementation and optimization of the modules to orchestrate access to and utilization of experimental resources;
- Task 1.3: Implementation and refinement of the modules for collecting, storing and visualizing measurement data.

As it was planned, Tasks 1.2 and 1.3 have followed a two-phase implementation, the first one being the development of a prototype implementation for the different system modules (completed during the first year). The second phase, the final complete implementation of the system, has been finalized at the end of the second year. Future hardware evolution of the platform or software maintenance might require further software updates. This will be carried out in the frame of the maintenance work package (WP4).

The main result of WP1 during the first year was a working prototype of the platform. This prototype has been deployed into production and is currently being used with real nodes running experiments and collecting measurements. We have collected extremely useful information from these early experiments, and a refinement of the design of the overall architecture is in place based on the feedback received not only from experimenters of the first open call, but also from MONROE developers during the deployment phase and when designing and testing EaaS solutions.

MONROE is compliant with the Fed4FIRE architecture, since the testbed entry point is an Aggregate Manager that implements the GENI Aggregate Manager API Version 3. MONROE is a registered project in the Fed4FIRE federation, and it reuses the same certificates issued by Fed4FIRE, and so, any already registered user in Fed4FIRE could be granted access to MONROE. However, the MONROE access portal and scheduler go beyond Fed4FIRE and provide REST-based APIs to enrich the expressiveness of virtualization and scheduling of nodes in the MONROE platform.

Furthermore, we have deployed a full integration with mPlane. MONROE platform gives native support for the mPlane protocol inside its own container, and the Tstat monitoring tool suite forms part of the collection of MONROE base measurements. The data collected by MONROE nodes is exported to MONROE database and can be seamlessly exported to any mPlane-compliant monitoring system.

The objectives of the Work Package 1 are included below, with a summary of the status of the completion of each objective, as of the second year of project activity.

- **Definition of MONROE use cases and analysis of system requirements (completed):** System requirements and use cases have been collected and documented in deliverable D1.1. A total number of 10 use cases, grouped in 3 broad categories (key mobile broadband metrics, application performance measurement, and innovative protocols and services), have been fully analyzed and documented.
- **Design a measurement system for mobile broadband technologies (3G/LTE and WiFi) (completed):** The overall architecture design of the MONROE platform has been completed (milestone MS1) and a working prototype implementation has been implemented and deployed (milestone MS4) in order to validate the seamless integration of system components and the fulfillment of use case requirements.
- **Design and implement Fed4FIRE-compliant modules for managing the access and the utilization of nodes (completed):** The MONROE scheduler system is compliant with GENI Aggregate Manager API Version 3, and the MONROE user authentication and authorization is ready to use certificates issued by Fed4FIRE certification authorities. The user management system and scheduling policies have been upgraded thanks to the feedback received from the initial users of the platform and in the process of defining EaaS and experiments carried out within the project by MONROE partners.

- **Design and implement the modules for collecting, storing, and accessing measurement data (completed):** MONROE prototype implementation provides a data storage and management system based on the combination of a big-data, NoSQL database and replicas for resiliency issues, a publisher/subscriber based solution metadata generation and for data transfer, and a visualization framework which has been enhanced and integrated with tools to visualize Tstat-generated data during the second year. We have added support for Fed4FIRE ETL middleware.

- **Implement an mPlane-compliant interface to allow the MONROE probes to be integrated into the mPlane architecture (completed):** The container-based architecture of MONROE node's middleware allows the deployment of multiple paradigms of measuring and data collection. A fully functional mPlane client is embedded in all MONROE's nodes, together with the Tstat monitoring tool.

All three tasks of this WP have been concluded and objectives have been reached by the end of the second year.

2.1 Progress in Task 1.1

The task was concluded at the end of the first year, as per the original schedule of Phase I activities.

2.2 Progress in Task 1.2

Task 1.2 comprises the implementation and refinement of the mechanisms needed for user access control and orchestration of the testbed measurements in presence of internal and external experimenters. Phase I was completed at the end of the first year with the release of the initial MONROE prototype. During the second year, with Phase II, we have worked at the final implementation of scheduler and user access/visualization portal using the feedback received both from the internal and external users. Based on the insights we gained during the prototype implementation and node deployment, we have modified and extended the system design and implementation with new functions, as follows.

- Refinement and improvement of the user access web system. The new user interface is more ergonomic and much richer (e.g., it includes resource type filters, resource state, and a map of resources that provides a quick-glance view of the whole system). It offers the possibility for users to specify additional run-time parameters for the Docker containers that are used to virtualize the MONROE nodes when it comes to provide node access to the experimenters. In this way, users can parameterize container behavior without modifying their images. The user interface server now uses valid SSL certificates, signed by Let's Encrypt. Additionally, user interface components have been upgraded to use only secure (i.e., "https") connections.

- Refinement and improvement of the experiment scheduling system, with access restrictions, user quotas and accounting. This required the design of data usage monitoring mechanisms directly reporting to the scheduler. A quota journal records all the transactions affecting user and node quotas. The users have access to usage reports and container logs in the user experiment results, which can be retrieved via the web user interface.

- Container certification. We have worked on the monitoring of deployment and analysis of container sizes, to facilitate the certification process of experiments. Furthermore, we have introduced the tools needed for the implementation and reporting of maintenance mode, preventing scheduling on nodes that fail the system self-test.

- Improvement of container configuration and management, including the implementation of a network name space for experiments. With the replacement of ZTE MiFis with Sierra PCI express modems at the end of the second year, the base system has provided a NAT layer between the modems and the MONROE network namespace. This has been implemented and tested based on veth interfaces, and an updated system configuration has been provided.
- Improvement of data collection management for experimenters. This includes the design and testing of a system to transfer the results of user experiments to a common repository accessible through the user client web.
- Production and continuous updating of a User Manual to guide external users through the full experiment workflow. Feedback from external users were very valuable in this activity.
- Creation of an SSH tunnel for containers, which allows external users of the platform to deploy their containers and connect interactively to them for debugging purposes. Access to the container is based on SSH access with public keys.
- Identification and inclusion of commonly used software packages in the base Docker container image provided to the external experimenters. This both reduces the download quota used by each user and the project as a whole.
- Public demonstrations of the platform at NetFutures, WoWMoM (awarded with the Best Demo Award) and Wintech/Mobicom.

With the submission of D1.3, this task has been concluded, and MS9 has been achieved in M24. All the updates to the User Access and Experiment Scheduling System will be carried out under the Maintenance work package, WP4.

2.3 Progress in Task 1.3

Task 1.3 concerns the design and implementation of the modules that collect the experiment results on nodes and servers, the database that will store the performance results of the experiments and the metadata, and the visualization tool that presents the results in a user-friendly public website. This task consisted of two phases. Phase I concluded with a prototype at the end of the first year. Phase II has been carried out and concluded in the second year, with the release of the final implementation with extended functionality and modifications.

The activities carried out to complete the tasks during Phase II are detailed in what follows.

- Continuous refinement of the metadata gathered by the nodes and saved to the MONROE database (e.g., modem battery level monitoring). This refinement consisted in (re-)designing the metadata format, investigating on key metrics to monitor and developing/testing the system for disseminating the metadata information in the nodes, developing of the backend system for storing and transferring the metadata and experiment data to the repository server and importing it into the database, and evolving the design of the database for storing metadata.
- Database architecture enhancement. The database schema has been improved to support the storage of metadata, serve the analysis of the planned use cases (traceroute, ping, http streaming, etc.) and allow for meaningful and efficient data extraction. We have also deployed a secondary copy of the MONROE Cassandra database. The two existing copies are located one in Sweden and one in Spain.

Access to the database replica is granted for all the external users based on SSH public keys. To reduce the risk of data corruption in our main database (Cassandra replicas are idem-potent), a mechanism to automatically add all the information from the previous day has been implemented. Therefore, the replica contains all the information from the main database with a delay of about 12 to 24 hours. Furthermore, to enable the foreseen integration with mPlane, we have designed new tables to integrate Tstat/mPlane data from the nodes in the MONROE database.

- Database data extraction. Daily CVS dumps of database activity are provided to users as a basic data access service (each sample only contains events for an interval of 24 hours). In general, the database can be accessed in two ways: building a custom application or using the standard cqlsh tool (e.g., to dump full tables). In addition, we have provided a simple application as a reference for users who want to extract data directly from the database. This application is essentially the same used to create the daily CSV dumps. Additionally, we have provided an application to dump the location of a node during a certain time period to a KML file, which can be later opened in tools such as Google Earth. We have also provided an additional example application that shows how to correlate GPS and MODEM (signal strength) information to generate KML files (for Google Earth) that show the modem coverage and link technology in different areas. All applications have been added to the project public repositories.
- Data validator. We have developed a data validator, which looks into the MONROE database to detect potentially faulty nodes and other error conditions. Further work will be devoted to this tool in the future—though in WP4, tackling maintenance issues—including a mechanism to automatically warn administrators of potential ongoing issues.
- Visualization tool enhancement and migration to a public space. The visualization tool has been enhanced and migrated, jointly with the database, to a new backend infrastructure to provide more redundancy and higher throughput. The visualization now includes results for http streaming experiments and Tstat statistics in addition to basic RTT (ping) and web download (http_download) data. To facilitate and optimize data views in the visualization tool (and in the user access interface), we have extended the scheduler REST APIs.
- Design refinement and testing of the user access system to setup, run, schedule, secure and collect experiment/log output from the experiment containers, using Fed4FIRE certificates. The repository mechanism for the results generated by external users has been implemented and is now accessible through the user web interface.

  With the submission of D1.3, this task has been concluded. All the updates to the Data Collection, Storage and Visualization System will be carried out under the maintenance work package, WP4.

3 Work Package 2

WP2 aims to manage and implement the procedures for setting up the MONROE experimental infrastructure in terms of hardware-related requirements. The WP2 work mainly consists of analyzing and selecting the hardware components, testing the materials in different borderline conditions and finally purchasing and assembling the MONROE nodes. Part of the activities carried out by the WP2 is also the distribution of the nodes to the partners in order to deploy the devices (both stationary and mobile) in Norway, Sweden, Italy and Spain. Moreover, WP2 is also devoted to implement the maintenance of the nodes with the development of proper recovery functions.

WP2 has been structured in three tasks:

- Task 2.1: Selection and Assembly of the Nodes (started on M1);
- Task 2.2: Deployment and Logistics (started on M7);
- Task 2.3: Robust Node Recovery (starts on M16).

The activities during the second year mostly focused on the assembly and deployment of the full set of stationary nodes, as well as of a part of the mobile ones in Task 2.2. Due to the issues we identified with some of the components (namely, the USB hub and the ZTE MiFi modems), the deployment of the mobile nodes has been delayed. A node re-design phase has been indeed needed to tackle these issues and propose a more robust and stable alternative solution. Task 2.3 has been successfully completed during this reporting period, with the design and implementation of robust node recovery procedures aiming to overcome specific node malfunctions. Task 2.1 was officially closed at the end of the first year, and only a small amount of appendix activities have been carried out to manage the acquisition of components for mobile and stationary nodes.

The objectives of the Work Package 2 are included below, with a summary of the status of the completion of each objective, as of the second year of project activity.

1. **Design of the hardware components of the MONROE nodes (completed):** After a thorough evaluation of the available alternatives in terms of suppliers, prices and delivery time, the hardware components of the nodes are selected. During the tests, the some of the components show clear problems. The consortium then reconsider the node design, and tested different components for the new node after extensive tests. The initial node design is documented in D2.1 and will be updated once the deployment is finalized in June 2017.
2. **Purchase of the selected components and assembly the pieces into MONROE nodes (ongoing):** Based on the initial node design, all the components are purchased and assembled. Due to the new node design, the consortium purchased the new components and is currently finalizing the assembly of nodes based on the final node design.
3. **Selection of the subscriptions and SIM cards for 3G/4G connectivity (completed):** The selection of three operators per country has been concluded. The agreement with the operators are negotiated and finalized. The SIM cards have been purchased and activated for the deployed nodes.
4. **Identification of the proper node hosts in which to deploy the whole infrastructure (completed):** Each partner has identified the proper hosts, especially the mobile hosts.
5. **Define robust node recovery functions (completed):** The consortium defined and implemented robust node recovery mechanism that discovers possible hardware failures and takes appropriate actions to solve the causes and restore the node's resources. The design and the implementation of the robust node recovery is reported in deliverable D2.3.

In the following, we provide a detailed per-task description of the activities carried out in WP2 throughout the second year of activity.

3.1 Progress in Task 2.1

Task 2.1 has the objective to select all the different MONROE hardware components building the nodes, their acquisition and their assembly. Basically all the activities were successfully carried out and finalized during the first year, with the submission of deliverable D2.1, which provided the MONROE hardware design choices, augmented with the definition of selection and testing processes, capabilities and limitations of the nodes.

However, towards the end of the second year, the consortium became aware of some obstacles while continuous monitoring the deployed node's status and performance. The major obstacles we experienced were:

- **Forced MiFi firmware updates:** In the last quarter of 2016, ZTE issued a forced update to the firmware of all MiFis. The update was applied even if the devices were configured not to receive automatic updates. After that, all our MiFis became inaccessible for the MONROE system. The possibility of new forced updates that can completely black out the MONROE platform is not acceptable.
- **Repeated Yepkit USB hub restarts:** During the first months of operation, we experienced repeated apparently random resets of the Yepkit hub. The problem was observed more frequently for the mobile nodes. The main reason for the restarts is due to the power variations in the input voltage of the hub which follows the specifications of the USB strictly. Especially for the mobile nodes, the converters might not provide such a stable voltage to the hub. Even more problematic than the restart of the hub is that after several of those events, the MiFis hang. In some cases, the only way to recover them is draining their batteries, or perform a (labour dependent) manual reboot by pushing the power button.
- **Unstable MiFis:** The modems themselves seem to be prone to resets or to enter a working state (transparent PPP) from which they can only be brought back into normal operation by draining their batteries or performing a manual reboot.
- **Bloated MiFi batteries:** After a few months of work, some of the MiFis showed clear signs of bloated batteries. This problem creates serious safety concerns for the deployment of nodes at places other than our own (controlled) premises.

Upgrade of the node design and implementation: The above described obstacles suggested the strong need to modify the node design, particularly to avoid clear safety risks. This then triggered a phase of evaluation for alternative node designs, taking into account to avoid MiFis and USB hub usage. We have evaluated in great detail different options.

Due to the safety risks, our main objective was to replace the MiFis with a battery-less modem while also solving the USB hub related issues. Regarding the modems, considering that anything with battery cannot be used in our system and USB modems are no longer produced, we chose to use miniPCIe modems. Among different miniPCIe modems, Sierra Wireless MC7455[8] is the only CAT6 modem available on the market (MiFis support CAT4), and it is proven to be quite stable in our tests. Therefore, the consortium decided to use MC7455 as the new modem in the MONROE nodes. To tackle the USB hub related issues, the consortium came up with 3 alternatives:

- **Dual-APU system:** This node design is based on a dual PC Engines APU2 system. One of the APU2 in each node has two Sierra Wireless MC7455 miniPCI express (USB 3.0) modems; the other has one MC7455 miniPCI express modem and a WiFi miniPCI express card. The advantage of this design is that it removes the USB hub and provides 3 miniPCI-e modems. The APUs are proven to work very stable and only one converter for the mobile nodes will be sufficient. One disadvantage of this design is that there is no space for an additional interface for the management traffic. Therefore, the management traffic should be managed together with the experiment traffic (e.g. management traffic should be scheduled).

[8] https://source.sierrawireless.com/devices/mc-series/mc7455/

- **Node with miniPCIe adaptor:** This node design uses a miniPCIe adaptor[9] that allows us to use Sierra MC7455 modems and connect them to USB ports. The adaptor boards do not implement protocol conversion, they just route signals from the USB pins of the miniPCIe slot to the modem board. These miniPCIe adaptors have external power supply, and this is necessary since APU cannot power both the internal slots as well as external USB ports simultaneously. With this node, we can use the 2 USB ports for MBB and use the two internal PCIe socket for the third MBB and Wifi. This node also removes the USB hub and MiFis, and similar to the first scenario, the management interface is removed.
- **Industry grade USB HUB with PCIe2USB modules:** This node design is closest to the current node with the necessary upgrades. First, it replaces the MiFis with miniPCIe modem modules[10]. This module has a SIM card reader and it is industry quality. Furthermore, it also relies on a industry grade externally powered USB hub such as Acroname[11]. This design needs minimum SW changes to the system, however, the new USB hub is extremely expensive and needs to be tested together with the miniPCIe modules.

We evaluated and extensively tested all the above three alternatives. We assembled and shipped test mobile nodes with different designs for a more accurate evaluation under mobility conditions. The main problem with the node design with miniPCIe adaptor was that during our tests, we experienced that the externally powered adaptor feeds power to the APU which causes the APU to crash after a certain period. The node design with the industry grade modem modules and USB hub was extremely expensive compared to the other two alternatives. Considering that APUs were quite stable and working well during the prototype implementation phase, we chose the dual-APU as our final node. Table 1 summarizes the evolution of the main platform from the prototype phase to the current state.

Table 1: Evolution of the main platform from the initial prototype to the implementation at M24.

Design aspect	**Prototype**	**Final release (New Node)**
Node platform	APU1D4	APU2
Node configuration	3xMiFis + WiFi	3xMC7455 + WiFi
Node HW	1xAPU + Hub + 3xMiFis + WiFi	2xAPU + 3xMC7455 + WiFi
Management Interface	Yes	No
Operating system	Debian 8 Jessie	Debian 9 Stretch
Modem Type	ZTE MF910 CAT4 USB MiFi	Sierra MC7455 CAT6 miniPCI express modem
Interface NAT	MiFi-provided	Node OS-provided

The new node design with dual-APU not only overcomes the safety risks and instabilities, but also presents additional advantages:

- We removed LiPo batteries from the node hardware implementation, eliminating any safety concerns raised by them and the burden of draining batteries after a modem crashes.

[9] https://techship.com/products/pci-express-mini-card-to-usb-adapter-with-external-voltage-and-casing/
[10] (https://techship.com/products/mpcie-usb-adapt-waterpr-casing/?signature=1036
[11] https://acroname.com/store/s77-usbhub-2x4?sku=S77-USBHUB-2X4

- We removed the USB hub, eliminating related instability issues and the single point of failure. With this, the node becomes more robust and we reduce the number of mechanical components that can fail during normal operations.
- When replacing the MiFi modems, LTE Category 6 ("CAT6") miniPCI express modems from Sierra Wireless became commercially available. The original Sierra Wireless modems used for management were LTE Category 3 ("CAT3") and the ZTE MiFis we used initially were only LTE Category 4 ("CAT4"). Therefore, replacing MiFi modems also allowed us to upgrade the MONROE platform to follow the latest industry deployments.
- The Sierra modems are simpler devices than the ZTE MiFis (which were full Android systems). Using them in the node implementation reduces overheads in the chain from the MONROE node protocol stacks to the MBB network.
- The prototype platform used the NAT mechanism provided by the ZTE MF910 MiFis. As the second iteration of the platform design uses only miniPCI express Sierra Wireless modems without NAT, this functionality has been moved to the operating system itself, where it can be controlled and configured by the MONROE platform.

Even though the task is officially closed, due to the new node design, an update to D2.1 is required. D2.1 will be re-submitted once the deployment is completed in June 2017.

3.2 Progress in Task 2.2

Task 2.2 aims at planning and managing the deployment of all the MONROE stationary and mobile nodes. This entails taking care of organizing logistics for shipment and placement of nodes, as well as for subscriptions for mobile operators SIMs across the different countries. During the first year, activities in Task 2.2 were mostly focused towards the definition of mobile nodes deployment strategies and all those logistic procedures needed to handle the collection of assembly material, nodes delivery in different countries, and subscriptions to different mobile operators in Italy, Spain, Norway and Sweden.

During the second year, some of these above activities have been refined and finalized, like the management of additional mobile operators subscriptions needed to cover newly assembled and ready to be deployed mobile and stationary nodes. However, the core part of Task 2.2 activities in this second year has been devoted to the deployment of the nodes according to the procedures and logistics defined in the first year.

In the case of stationary nodes, we successfully managed to finalize the delivery of all the planned 100 stationary nodes. In addition to deployment, each partner was also involved in testing activities over the stationary nodes. Moreover, with the aim of satisfying the needs of external users from the first open call, we successfully managed to assembly and ship around 20 additional stationary nodes to the external users. However, the assembly and deployment of mobile nodes has been delayed during the second year of the project. In total, we managed to assemble and ship around 70 mobile nodes, most of them shipped and deployed in buses, trains and trucks. In addition to some long delivery time for some of the components, the technical issues described in Section 3.1 have delayed the mobile nodes assembly and deployment procedures.

After extensive evaluation of different new node design alternatives, dual-APU system is chosen. This choice required an upgrade of already assembled (and in most cases) deployed mobile nodes, with additional hardware to be acquired for new dual-APU nodes. Indeed, we started at the end of the second year with the assembly of new dual-APU systems, upgrading (where possible) already available nodes. Though much more robust and stable, the new node implementation with dual APU is also more expensive, which impacts the total number of nodes we initially planned to integrate in the MONROE testbed. Currently, the distribution of nodes per country with MONROE coverage is the following:

- Norway: 35 deployed nodes (25 mobile and 10 stationary), 5 testing nodes
- Sweden: 45 deployed nodes (30 mobile, 15 stationary), 5 testing nodes
- Italy: 45 deployed nodes (25 in public transport vehicles, 15 in delivery trucks, 5 stationary), 5 testing nodes
- Spain: 10 deployed nodes (all stationary), 5 testing nodes

This adds up to a total of 135 deployed nodes and 20 testing nodes. We mention that we are currently in the process of returning the ZTE MiFis we have purchased to their manufacturer. Depending of the outcome of this process (if/when we receive the money back), we will use that money to build more dual-APU2 nodes, hence increasing the number of nodes in the MONROE testbed.

We note that several of the external users also bring extensions to the MONROE platform by purchasing full operational nodes (and funding the corresponding data subscriptions) and hosting them in their premises. This will allow us to increase the total number of nodes in the testbed and diversify the geographical coverage of the platform.

As an intermediate outcome of assembly and deployment activities, deliverable D2.2 was released in September 2016 and provided the specification of logistics workflows (material acquisition, delivery, etc.), nodes assembly procedures and testing methodologies, and strategies for efficient and effective deployment of nodes.

Even though the task is officially closed, the plan for the beginning of the third year is to continue with transition from the old hardware system to the new dual-APU system. Due to the new node design and delays to the deployment, an update to D2.2 is required. D2.2 will be re-submitted once the deployment is completed in June 2017. The third year efforts will be accounted as part of the maintenance activities in WP4.

Progress in Task 2.3

Task 2.3 main objective is to implement a robust recovery procedure to prevent unexpected malfunctions, which may run the nodes into an unstable condition. The start date of this task was originally planned on M16, however, the consortium agreed to advance in this task earlier, during the first year, considering the importance of its outcomes in terms of stability and sustainability of the platform. Main outcome from the first year was the complete identification and definition of workflows for recovery actions and events, which put the basis for the development activities.

During this second year, activities in Task 2.3 have been focused to the implementation of robust node recovery and associated maintenance methods within the MONROE platform. These methods are key for an effective and continuous operation of the platform, and to detect problems early in advance to reduce manual intervention to a minimum while ensuring that failed hardware components are timely identified and replaced. The node recovery prototype provides the fundamentals for the maintenance routines and maintenance activities carried out in WP4. From a technical perspective, dedicated mechanisms to monitor the operational status of the deployed nodes have been implemented in the form of a set of watchdog functionalities, which are able to detect problems in a timely fashion and take appropriate corrective actions. The core part of the robust node recovery prototype enables hard restart, bringing the system back to a known working baseline without manual intervention. Dedicated procedures for handling node failures due to hardware malfunction have been implemented, including malfunction detection, cause identification and proper recovery function enforcing.

Task 2.3 is concluded during the second year, and all the above different node recovery activities have been reported in D2.3. Any needed updates at the robust node recovery will be carried out in maintenance workpackage, WP4.

4 Work Package 3

WP3 deals with the experiments that we aim to accommodate in the MONROE system. It involves the design, implementation and execution of experiments matching different use-cases for the MONROE platform, as well as providing the experiments as a service (EaaS) to external users. WP3 integrates two main tasks:

- Task 3.1: Design and run experiments for each use case;
- Task 3.2: Provide the designed experiments as EaaS.

The activities during the second year mostly focused on the design of the experiments covering the main use-cases of the platform. All the designed experiments are further provided as EaaS to the users of the platform throughout the second year.

The objectives of the Work Package 3 are included below, with a summary of the status of the completion of each objective, as of the second year of the project activity.

1. **Design and execution of experiments for each use case defined in WP1 (ongoing):** The consortium has designed and executed various experiments covering all the use cases during the second year. During the last year of the project, the consortium will expand the range of experiments.

2. **Assessment and monitoring of MBB networks (ongoing):** The consortium has started analyzing the results of experiments to better understand the mobile networks in the wild. Due to the hardware problems, large scale monitoring of the mobile networks has been delayed but the consortium will spend significant resources on assessment and monitoring of the MBB networks in the last year of the project.

3. **Identification of the key performance metrics of MBB networks that affect user experience (ongoing):** The consortium specifically focused on coverage analysis and speedtest like measurements. We have run initial tests verifying our experiment design. Currently, we are in the process of collecting large scale measurements and analyzing the results. With the large scale platform in place, the consortium is going to accelerate these efforts in the last year of the project.

4. **Test and evaluation of innovative services and applications for MBB networks (ongoing):** The consortium ran experiments focusing on application performance, more specifically, performance of web and video application. To this end, we collected measurement results mostly focusing on the stationary nodes. The focus will be shifted to the analysis under mobility in the third year. Regarding the innovative protocols, the consortium investigated different aspects of mobile networks in order to understand the mobile ecosystem. These works cover the path support for different protocols and impact of different protocols on the performance. In the third year, we will extend these measurements while also focusing more on innovative protocols such as MPTCP.

5. **Provision of the designed experiments as a service to extend the width and breadth of the platform's application range (ongoing):** The consortium has provided all the designed experiments as EaaS in the project's public Git repository. Moreover, the external users contributed to this effort by sharing their codes with other users. In the third year, the consortium will continue providing all designed experiments as EaaS and will encourage all the users of the platforms to do so.

4.1 Progress in Task 3.1

Task 3.1 involves design and running of experiments for use cases defined in WP1.

We organized the MONROE use cases are organized under three main categories: (i) key mobile broadband metrics, (ii) application performance and (iii) service and protocol innovations. For each category, we have carried out numerous discussions with academic and industry partners in order to identify meaningful examples of experiments. We have further conducted a thorough analysis of the state of the art on MBB measurement and assessment techniques. These efforts resulted in a set of experiments to be implemented by the consortium in order to enable each use case.

We next detail the experiments have been designed and developed based on the three use cases defined in WP1.

Key MBB parameters: The core functionality of the MONROE platform is to provide experimenters a rich dataset of key mobile broadband metrics, from which different stakeholders can further extract the information of interest regarding the performance and reliability of MBB networks. To measure the network in a reliable and fair way, we identified several metrics that (accurately) capture the performance and the conditions under which we evaluate these metrics. We implemented several experiments to monitor these key metrics, as follows.

- We developed a ping (RTT) container and HTTP downloading experiment (throughput) container, which are running as part of the continuous base experiments. With this, we ran an initial measurement campaign on monitoring and assessment of mobile networks. We used this initial measurements for the submission of a joint consortium paper on the MONROE platform to IEEE Communication Magazines (rejected from publication, will be resubmitted). Throughout the second year, we brought several modifications to these implementations. In light of the node implementation upgrade, we performed minor modifications to align with new design requirements. Notably, we enabled the HTTP experiments to run on a wider set of setups (e.g., with self signed certificates, on different ports etc).

  Meanwhile, we have also implemented the containers for UDP and TCP ping experiments, which will be used in conjunction with the current ICMP ping experiment.

- We have implemented and tested containers that perform traceroute and paris-traceroute probing. They can be scheduled as an internal experiment or used directly by any experimenters. A binary version of paris-traceroute is included in the base image, instead of the repository version that is problematic in the testbed. With this binary version, an example container will be written to take parameters from the scheduler and run paris-traceroute experiments. An effort has been made to reduce the potential impact on node availability for user experiments. Evaluation was done towards the feasibility of launching several traceroute experiments in parallel with raw sockets.

- We have been maintaining the container running continuous Tstat monitoring. The Tstat installation process has been finalized and currently all MONROE nodes deploy Tstat. We have been actively exploiting the Tstat rich output, by analyzing the log files we collected during HTTP experiments, together with the metadata collected by the MONROE nodes.

- We ran several bulk upload experiments on mobile nodes over TCP. A server has been installed and mobile nodes run experiments based on an iperf container. Traces were collected on both MONROE nodes and backend server. This container was tested and scheduled on the nodes in different countries.
- We leveraged access to the rich modem metadata to build coverage maps of different operators MONROE measures. We used this information to further characterize and then identify the coverage profile of MBB networks. Our paper on profiling MBB coverage has been accepted and received the best paper award in IFIP TMA workshop. An extended version of the paper has been invited to the Elsevier Computer Networks special issue on Machine learning, Data Mining and Big Data frameworks for Network Monitoring and Troubleshooting.
- We have implemented a speedtest experiment that is in line with the current test regulators use in different countries. We are currently in the process of testing the corresponding container and we are planning a later large-scale deployment. We have surveyed different approaches for speed measurements (from academia or industry) and aim to use MONROE in order to compare their performance and accuracy.

Application performance: The experiments we implemented in this category focused on web and video services, such as video streaming, video surveillance and web performance.

- We have finalized the implementation of a container that is able to run experiments on web performance. The container design and implementation relies on a selenium web driver based tool (Headless Browser) in python to fetch web resources and measure the web performance in a headless browser using Firefox with different HTTP protocols. The container "HeadlessBrowser" uses selenium as a web application testing framework. This container was used to measure the performance of different HTTP protocols (HTTP/1.1, HTTP1.1/TLS, HTTP2) for popular sites. An analysis started on the impact of different protocols on the objective QoE for the end users. We are currently working on a webperf tool that is able to quickly analyze the performance of web services based on the output of the HeadlessBrowser. Preliminary analysis has started to investigate how the web service is impacted by the underlying channel conditions (*e.g.*, by analyzing the HTTP traces as well as headless browser measurements together with metadata).
- We have been actively working on the design and implementation of experiments on adaptive video streaming with MPEG-DASH. We have implemented and tested a DASH client called AStream. Initial performance measurements, captured from the stationary nodes have been analyzed whereas the traces of DASH videos from mobile nodes are stored. In order to move closer to end users, quality of experience (QoE)-related experiments were developed and used on the user devices. Two mobile QoE apps were implemented. The first one is associated with AStream that downloads DASH video files from the server, plays back at user devices and gathers user experience feedback in real-time. The second mobile app directly emulates DASH on mobile devices and collects DASH performance data from Android libraries while gathering user feedback. Data analytics is used to correlate the collected performance metrics from various sources and then create a QoS-QoE mapping for video streaming services, followed by the establishment of a QoE model for MPEG-DASH services.

- Using the existing data MONROE generates, we have investigated video surveillance-related questions. Two major problems were considered: i) optimize the cost of uploading videos from security cameras installed on buses; and ii) video upload scheduling problem. Optimal solutions have been found for the cost optimization problem. Three heuristics have been proposed and primiarily tested using trace driven simulations. Results have been presented at the *SmartCity16: The 2nd IEEE INFOCOM Workshop on Smart Cities and Urban Computing* and at the *INFOCOM Student Activities workshop, 2016*. For the upload scheduling problem, an adaptive scheduler was designed and evaluated for deadline-constrained content upload from Mobile Multihomed vehicles.

- Through the projects we have funded after the two MONROE open calls for External Experimenters, we expect to add other experiments that enable the evaluation of different applications in MBB (see Section 6 for a full list of experiments from external users). A notable example is the YoMoApp (YouTube Performance Monitoring Application) tool that enables performance analysis of mobile networks with respect to YouTube traffic, and developing optimization solutions and QoE models for mobile HTTP adaptive streaming. The tool is already running in the MONROE platform and is also available as an EaaS to other potential MONROE users.

Innovative protocols: MONROE not only provides measurement for existing protocols and services, but also allows for the flexibility of testing and assessing innovative protocols and services.

- Together with partners from the H2020 MAMI project, we have been working on supporting in MONROE a tool able to test path support for IP and TCP functionality (*e.g.*, MPTCP, ECN, TCP Fast Open). To this end, we customize PATHspider[12] to enable measurements across mobile broadband paths in MONROE[13]. PATHspider is a tool developed for A/B testing of path transparency. It allows to test the feasibility of deploying new protocols in the Internet and quantify the impact of path impairments. We submitted a paper describing the customization of PATHspider for MONROE, together with proof-of-concept measurements in MONROE to the IEEE/IFIP Workshop on Mobile Network Measurement (MNM17) held in conjunction with the Network Traffic Measurement and Analysis Conference (TMA).

- We have also been focusing on measurements for path support for generalized ECN[14], which aim to assess the feasibility of adding Explicit Congestion Notification to TCP control packets and TCP retransmissions. To this end, we use `tracebox`[15] to discover the hops along a path that may be bleaching the ECN code points.

- Modern networks often rely on dedicated hardware components generically dubbed middleboxes to perform advanced processing functions like, for example, enhancing application performance (e.g., traffic accelerators, caches, proxies), traffic shaping (e.g., load balancers), optimizing the usage of IPv4 address space (e.g., NATs) or security (e.g., firewalls). One major issue arising from this approach is that middleboxes might, in some cases, filter traffic that does not conform to expected behaviors, thus ossifying the Internet and rendering it as a hostile environment for innovation. Using measurements with Tstat against a server under our control, we have worked towards identifying the presence of proxies in mobile broadband operators' infrastructure. We submitted the results of this work to the International Teletraffic Conference (ITC) in March 2017.

[12] https://pathspider.net/
[13] https://github.com/mami-project/pathspider-monroe
[14] https://tools.ietf.org/html/draft-bagnulo-tcpm-generalized-ecn-02
[15] http://www.tracebox.org/

4.2 Progress in Task 3.2

This task focuses on delivering a set of selected experiments to external users as EaaS and extending MONROE's application coverage. It aims to achieve the following goals:

1. Provide detailed experimental information such as scripts, user instruction/manual, and short description of expected results.

2. Simplify and speed up the learning process for external users to use the MONROE platform and start experimenting quickly.

3. Allow external users to modify the experimental scripts based on their demands, in compliance with MONROE's objectives.

Throughout Year 2, consortium partners have continuously worked on enabling selected experiments usable by external users. All designed experiments are available as pre-built containers that can be directly scheduled by users, and also available in source format published at GitHub (https://github.com/MONROE-PROJECT/Experiments/tree/master/experiments). The sources can be used to build and run the experiments as is, extended, or used as a template for other custom experiments.

- A template experiment container has been developed and later adapted to support new platform functionality, *e.g.*, internal container interface naming, code to handle modems disappearing/loosing connection, new metadata, etc. To ease the adaptation of the platform, the template container is split into two: A simple "Hello world" that shows the system and how one can use the scheduler, retrieve data, etc.; and a more complex template that illustrates a container that reacts on metadata and handles modem events.

- The two traceroute containers (traceroute and paris-traceroute) have been uploaded to the project repository in github. They can be directly employed by users in the platform.

- The ping and HTTP download containers are provided to external users with updated documentation and code to reflect the current node design and to ease user adaption. Adaptations have been made to improve data accuracy and fulfill new requirements based on the lessons learned while running the platform.

- Access to Tstat passive traces is provided for the experimenters to allow them correlate Tstat traces with their experiment data.

- We have implemented an AStream container, which is a very light weight DASH video client. The AStream container can directly be scheduled on the nodes or modified based on experiment needs. Documentation and code for the Astream container has been updated to accommodate user's requests.

At the end of the second year, D3.1 has been prepared and delivered. That document includes all the information relative to the experiments that are currently available for use by experimenters, including detailed instructions on their usage. That information has also been included in the user manual to achieve a broader dissemination.

5 Work Package 4

Maintenance of the testbed infrastructure is a crucial task. Clearly, experimenters making use of the testbed depend on its availability for their experiments. The main focus of this WP to provide maintenance of the testbed. WP4 integrates three main tasks:

- Task 4.1: Routines for Maintenance
- Task 4.2: Node Maintenance
- Task 4.3: Backend and Data Maintenance

The maintenance WP started in M16. During the second year, the consortium first identify and deploy maintenance procedures (Task 4.1) to keep the testbed infrastructure up and running. Then the focus was shifted to the main maintenance activities: Node maintenance including troubleshooting of software-related issues as well as hardware failures (Task 4.2) and Maintenance of backend infrastructure together with maintenance of research data produced by experiments (Task 4.3).

The objectives of the Work Package 4 are included below, with a summary of the status of the completion of each objective, as of the second year of the project activity.

- **Keep the nodes in operating condition (ongoing):** The consortium defined maintenance procedures to keep the nodes alive. Besides automated monitoring tools and ticketing system, the regular weekly telcos brought the maintenance team together. This way, we were able to react to the issues fast and keep the system up and running. We will continue helding the maintenance telcos to prioritize the uptime of a large number of nodes during the third year of the project.
- **Keep the number of nodes that are operating large enough at all times (ongoing):** During the second year, due to hardware related issues discussed in Section 3.1, the number of nodes that are operating was around 50. The platform is now going through a transition to the new node design and will reach a large number of nodes in June 2017 (155 nodes provided by the consortium and 20+ nodes provided by the external users). Our main goal for the third year to ensure the high number of nodes operational.
- **Keep the new research data obtained from experiments on nodes in the backend accessible to the experimenters (ongoing):** During the second year, the consortium took necessary actions to have backups of the data produced by the experiments. We also provided the replica of the database as well as the daily dumps of the csv files. We will continue these activities during the last year of the project.

5.1 Progress in Task 4.1

Task 4.1 involved coming up with routines for maintenance. Towards this end, the consortium first identified the maintenance goals as: keep alive, react fast and data first. The consortium then defined the entities to be monitored and how they are monitored as well as the maintenance procedures. In order to run the maintenance activities efficiently and to react timely, we aim to automate as many tasks as possible. However, not all issues can be resolved automated. For all other maintenance activities, we created a Maintenance Ticketing system where we follow maintenance issues with an issue tracker, log maintenance actions and distill repeating actions into maintenance routines. To be able to address the issues that require manual intervention, the maintenance team comes together in a weekly telco to solve the issues collectively.

This task is concluded with the submission of the report on maintenance routines in Deliverable D4.1 and Milestone 7 is achieved.

5.2 Progress in Task 4.2

Task 4.2 Node Maintenance (HW+SW) involved both regular maintenance operations. For the regular operations, one partner was selected per country to provide the maintenance role:.

- Italy: NXW monitored the status of the nodes in Italy for the stationary nodes in Pisa and in Torino (in collaboration with POLITO). The mobile nodes have been also verified periodically. Manual interventions on the stationary nodes in Pisa have been done. Some of these nodes (both stationary and mobile) have been switched-off or disabled due to the problem of inflating ZTEs. POLITO has worked on the verification of the correct status of the stationary nodes in Torino. Some problems with the SIM cards and the configuration of the nodes have been reported and solved.
- Sweden: KAU has periodically verified the nodes in Sweden. Manual actions have been taken to reboot the nodes or switch on/off the MiFis and do further troubleshooting in cases where this action has not helped to have the node fully functional.
- Spain: IMDEA monitored the nodes in Spain to verify their correct functioning and help to improve overall system and communication stability. Manual interventions have been performed to debug some issues.
- Norway: SRL has been monitoring the nodes in Norway and was involved in troubleshooting of node HW and SW for all the deployed nodes. Moreover, some hardware repairs on the nodes has been carried out for the nodes placed in Norway.

The root cause analysis of any newly discovered issues, implementation of repair actions and adaptation of maintenance routines has been performed, as well as regular and administrative maintenance action on all nodes. Furthermore, we worked on evaluating alternatives for the new design based on APU2 boards. The new design is mostly motivated by the issues experienced with MiFis and USB Hub. Regarding the MiFis themselves, they were spontaneously updated and had passwords change, leaving the modems unaccessible until the MiFi-unique passwords could be figured out. Further, issues with battery bloat in the MiFis were investigated. As part of the trouble shooting, it was also found that the SIM power sockets were a cause of significant power transients in periodic intervals. With regards to USB, the Yepkit hub proved to be unstable and caused disconnections and bus resets, and we have tested other USB-hubs, without hub, with other modems (Sierra modems in enclosures), as well as attaching fewer MiFis to reduce power requirements. We updated the MONROE operating system with a more stable configuration management, and prepared an update of the operating system version from Debian jessie to Debian stretch. We also updated the management inventory to a new version with extended features, showing node event history and the option to show traffic reports.

5.3 Progress in Task 4.3

The objective of this task is to provide troubleshooting of backend and database access issues throughout the project.

We have worked on the backup mechanism for the log files received from the nodes. External users can now access a MONROE database replica with near-live (12-24 hours delay) data. We have also completed the development of a mechanism to produce daily dumps of the MONROE database in CSV format. These files are available to external users through the user interface web server. We have further focused on the data management to ensure that all the nodes report proper information to the database servers.

We have modified the architecture for transferring the internal experiment results. The restructured system result in fewer and more efficiently compressed files which both reduces the quota usage during transmission and efficiency of the database import itself. We have further monitored the backend systems and servers to ensure their correct work.

6 Work Package 5

WP5 concerns the support for external users of the MONROE platform. This WP started in M10 and will last until the end of the project. Making mobile broadband experimental facilities available for experiments of novel protocols and services is a key target for MONROE. To succeed in this, the experimentation and feedback provided by external users is highly important and half of the budget in MONROE have been allocated for external experiments. Two open calls for external users have been arranged within MONROE: the first call has specifically targeted researchers, innovators and businesses that depend on MBB networks. The second call was open to all user groups. This work package covers all the activities related to these calls and the support for external users. It is organized into two tasks, where the calls for users and user selection is handled within Task 5.1 and the required support to the external users is provided within Task 5.2.

The objectives of the Work Package 5 are included below, with a summary of the status of the completion of each objective, as of the second year of the project activity.

- **Make the MONROE experimental platform available to external users through open calls (completed):** The consortium selected 27 external users to join the project through two successful open calls.
- **Provide efficient support for external users of the MONROE platform (ongoing):** The consortium assigned a patron to each open call user to ensure providing efficient support to the users. Furthermore, the consortium organized three successful workshops with the users and implemented a mailing list to have a forum for platform related discussions among all the users. In the third year of the project, all the patrons will continue working closely with their users and the consortium is planning to hold at least 2 more user workshops to discuss the MONROE experience and the experimental results.
- **Gain input on how to extend and enhance the MONROE platform from external users (completed):** During the second year, the consortium collected feedback from the first open call users on how to improve the platform. This feedback is reported in D5.1 and is taken into account for the final system implementation.

6.1 Progress in Task 5.1

The consortium publicly announced the first open call during Y1 (on December 23, 2015), but the selection of the call one users was carried out during Y2. The first open call had a total budget of 1,8 million EUR, and targeted 12 proposals (maximum funding per proposal of 150.000 EUR) for funding. Deadline for proposals submission was set to M13, mid-March, and 40 proposals were received.

The main work by the consortium in Y2 was to organize the review process to evaluate the applications and select its first set of external users. A call committee consisting of Anna Brunstrom (KAU, chair), Ozgu Alay (SRL), Hakon Lonsethagen (Telenor), Pedro Andres Aranda Gutierrez (Telefonica, external representative) was appointed. A total of 26 reviewers, recruited both externally and from within each of the organizations participating in MONROE, delivered over 120 reviews to provide at least three independent reviews for each proposal.

Each proposal was reviewed by reviewers from both academia and industry and by at least one external reviewer[16]. The proposals that were ranked above the threshold, had an average score of 10 or higher, were

[16] There was one exception to this where the external reviewer could not deliver his review. As the proposal was already poorly ranked and to not delay the review process, the third review for this proposal was also provided by a reviewer from one of the MONROE organizations.

then checked for compliance with the MONROE requirements and feasibility with the platform by Anna Brunstrom (KAU), Ozgu Alay (SRL) and Vincenzo Mancuso (IMDEA).

All proposals were then discussed in a meeting with the call committee. The seven highest ranked proposals, with a score of 12, were all selected. The committee then selected an additional five proposals amongst the 14 proposals that were given a score of 11 by the reviewers. In line with the call announcement, the call committee prioritized having a diverse set of proposals that provide feedback on the full potential of the MONROE platform in their selection.

The following 12 proposals were selected:

- PoWeR: Performance of Web RTC-enabled services on the MONROE platform, Industrial Innovation, N. Amram Technologies Ltd, Israel.
- Napplytics: Mobile Network Analytics for Apps Performance Design, Scientific Excellence, EUROB CREATIVE and UNIVERSIDAD POLITECNICA DE VALENCIA, Spain.
- Nestor: Quality-of-Experience of adaptive video streaming in mobile broadband networks, Industrial Innovation, StreamOwl (SO) and Athens University of Economics and Business (AUEB), Greece.
- Sometime: Software defined network based available bandwidth measurement in MONROE, Scientific Excellence, University of Napoli Federico II, Italy.
- Movement: Extending and Experimenting upon the MONROE platform towards Voice and VidEo StreaMing AssEssmeNT in Mobile Networks, Industrial Innovation, COSMOTE Mobile Telecommunications S.A. and FERON TECHNOLOGIES P.C. Greece.
- Monroe-LTE: AFFORDABLE LTE NETWORK BENCHMARKING BASED ON MONROE Acronym, Industrial Innovation, Allbesmart LDA, Portugal.
- Sophia: Software Radio for Measuring Mobile Broadband Networks, Scientific Excellence, Software Radio Systems, Ireland.
- Maril-in-Monroe: Measurement Adaptation and Reporting in LTE, Scientific Excellence, University of the Basque Country, Spain.
- Ricercando: Rapid Interpretation and Cross-Experiment Root-Cause Analysis in Network Data with Orange, Scientific Excellence, University of Ljubljana, Slovenia.
- PREC: Prioritisation and Resilience for Emergency Communications, Scientific Excellence, University of Aberdeen Court, Scotland.
- Mobi-QoE: Monitoring and Analysis of Quality of Experience in Mobile Broadband Networks, Scientific Excellence, AIT Austrian Institute of Technology GmbH Tran-Gia Informationstechnik GmbH, Austria.
- Unique: Utility-based Networking experiments for Improving QUality of Experience in mobile broadband environments, Scientific Excellence, Institute of Communications & Computer Systems (ICCS) and Incelligent (Incel), Greece.

The MONROE second open call was publicly announced in late September, 2016 via the project website, twitter and by sending the news to the list of subscribers interested in knowing more about the project. Various mailing lists and other FIRE related channels were also used for the announcement. We prepared the open call, putting together the detailed guide for applicants and the proposal template. The call announcement was similar to the announcement for call one, but with a further emphasis on the sustainability of the

MONROE platform. The second open call had a total budget of 1,5 million EUR, targeting 15 proposals (maximum funding per proposal of 100.000 EUR) for funding. Deadline for proposals submission was December 2, 2016. A total of 50 applications were received.

Following the submission deadline, the review process to evaluate the applications and select a second set of external users was organized by the consortium. The call committee remained the same as for the first open call. A total of 32 reviewers, recruited both externally and from within each of the organizations participating in MONROE, delivered over 150 reviews to provide at least three independent reviews for each proposal.

Similar as for the first open call, each proposal was reviewed by reviewers from both academia and industry and by at least one external reviewer. The top 23 proposals that all had an average score above 10 were then checked for compliance with the MONROE requirements and feasibility with the platform by Anna Brunstrom (KAU), Ozgu Alay (SRL) and Vincenzo Mancuso (IMDEA).

All proposals were then discussed in a meeting with the call committee. The 13 highest ranked proposals, with a score of 11.5 or above, were all selected. The committee then selected an additional two proposals amongst the four proposals that were given an average score of 11.3 by the reviewers. In line with the call announcement, the call committee prioritized platform sustainability and having a diverse set of proposals in their selection.

We selected the following 15 proposals:

- Characterizing Carrier Grade NATs in Mobile Broadband Networks: CGNWatcher, Scientific Excellence, Universidad Carlos III de Madrid, Spain.
- Network Neutrality in Mobile Broadband: NeutMon, Scientific Excellence, University of Pisa, Italy.
- Reconstruction of operator policies in MBB networks for improved user experience: RECON, Scientific Excellence, Eötvös Loránd University, Hungary.
- Experimental validation of REM-based machine learning algorithms for SON using MONROE nodes: MONROE-SON, Industrial Innovation, RED Technologies SAS, Instituto Politécnico de Castelo Branco, Portugal.
- Multi-homing with Ephemeral Clouds on the Move: MEC, Scientific Excellence, University of Macedonia, Greece.
- Programmable and Robust Smart Grid Data and Control: RASnet, Scientific Excellence, DAI-Labor, TU Berlin, Germany.
- Towards end-to-end Multipath TCP, Scientific Excellence, Université catholique de Louvain, Belgium.
- Smart City Security Monitoring Platform: Cloud Eyes, Industrial Innovation, Institute of Bioorganic Chemistry of the Polish Academy of Sciences, Poznan Supercomputing and Networking Center (PSNC), and TechInnowacje sp. z o.o., Poland.
- Optimization of QoE of Mobile Broadband Services through Machine Learning: OPTIMAL, Industrial Innovation, Modio Computing SP & Paris Descartes University (PDU), France.
- Fast and Lightweight Capacity Benchmarking of Mobile Broadband Networks in MONROE: FaLiCaB, Scientific Excellence, TU WIEN / Institute of Telecommunications (TUW), A1 Telekom Austria AG (A1), Austria.

- Network Self-Optimization based on End-To-End measurements: eSON, Industrial Innovation, Universidad de Málaga (UMA), Spain.
- Dynamic Pricing in HetNets: DAPHNE, Scientific Excellence, University of Thessaly (UTH), Greece.
- Traffic and Data Offloading in Mobile Networks: TTOff, Scientific Excellence, University Politehnica of Bucharest, Romania.
- Characterising Mobile Content Networks in the Wild: CaMCoW, Scientific Excellence, Queen Mary University of London (QMUL), United Kingdom.
- Feasibility study of latency-critical connected vehicle applications in MONROE: FELICIA, Scientific Excellence, SICS Swedish ICT, Sweden.

In connection to the open calls, we have prepared the contract to be signed with external users in order to make them Associated Partners to the consortium. Prior to the call deadlines, we also handled all the queries from external users regarding the call and the proposal submission for both open call one and open call two.

With the successful selection of the second set of external users this task has been completed. The procedure established for announcing the open calls and the set up of the evaluation process is documented as part of D5.1. With the selection of the phase I and phase II external users MS6 and MS11 are achieved.

6.2 Progress in Task 5.2

This task started in M13 and aims to provide adequate user support for the external user to allow them to work on their experiments and HW/SW extensions as efficiently as possible. To support the users a patron has been assigned for each external user. For the users from the first open call IMDEA serves as the patron for five users, KAU for four users, SRL for two users and POLITO for one user. Throughout Y2, all patrons have held meetings with the projects they patron to follow up on their progress and have collaborated to provide user support. In particular, we have provided support for using the certificates for accessing the platform and scheduling experiments, improving the system to better suit the needs of the external users.

During early Y2, we prepared a first version of the User Manual. The User Manual introduces the capabilities of the platform and explains how a user gains access to the system and executes his/her experiments. Later, we distributed the updated User Manual to the external users and also made it publicly available. Following public release, the content of the User Manual has been continuously updated to capture the continous development of the MONROE platform.

Two User Workshops have been arranged during Y2 to introduce the call one users to the platform and follow up on their progress. The first User Workshop for experimenters was held in Oslo, June 13-14, 2016 at Simula Research Laboratory.

To gather input on the experiences with the MONROE platform and get suggestions for improvements, all external users were requested to provide a First Feedback Report to the consortium. The Feedback Reports were delivered in the end of September and served as an important starting point for the second User Workshop.

The second User Workshop for experimenters was held in Madrid, October 24–26, 2016 at the IMDEA Networks Institute premises. During the workshop, the consortium analyzed the feedback received from the external users, which was generally positive. Several issues on user interface usability, scheduling policies and experiment design were analyzed and addressed. Some of the main outputs of the sessions were the agreement by the consortium to include additional functionality in the MONROE base container images, such as Firefox support. Coding sessions were also conducted to address specific issues encountered by the different projects, such as experiment design, real-time metadata access during experiment execution on the nodes and database access for off-line analysis.

Overall, user support during Y2 has worked very well and the open call one users have started to form an active MONROE user community. During Y3 the open call two users will be welcomed into the MONROE project, and planning for their support has also been initiated. Patrons for the call two users have been assigned and planning for the third User Workshop in Pisa at the end of March 2017 has been initiated.

7 Work Package 6

The description of the work carried out in WP6 (Management) is structured around the main objectives of this WP, as follows.

- **Conduct and control external communication by the project management to EC:** During the first two years of the project, SRL communicated with EC on behalf of the consortium. The deliverables are submitted on time. The changes in the consortium were communicated to EC and SRL has submitted an amendment due to departure of NET1, reallocation of subscription costs and overall extension of the project.
  - Departure of NET1, reallocation of subscription costs and project extension with by five month: NET1 has left MONROE's consortium due to lack of resources after their representative in MONROE has resigned. Consortium decided not to find a substitute beneficiary for NET1, therefore, NET1's original budget and tasks were redistributed between KAU and SRL. SRL is now the lead partner for D7.3 - Exploitation on Sustainability. NET1 had a total of 8 man-months in the budget; 2 for WP1, 2 for WP2 and 4 for WP7. Two person-months from WP2 were transferred to KaU, who took over the deployment and logistics of the nodes in Sweden. The budget for the other 6 person-months were transferred to SRL under the category other direct costs. This represents extra 67745 Euros for SRL. The reason for transferring part of NET1's budget to Simula other direct costs was that NET1's main contribution to MONROE would have consisted in hardware (150 modems for the nodes) and 100 high quota subscriptions in Sweden as in-kind. Total estimated budget for the consortium remains the same. In January 2017, the consortium submitted the amendment to the European Commission. The main change to the original MONROE proposal and the last valid amendment from March 2015 is the removal of NET1 from the consortium. Since there is no partner replacement, the budget and the tasks were redistributed among KAU and SRL. In addition, redistributing the subscription costs simplified the administration process for the partners. It further addressed the necessary budget and classification adjustments in order to allow hardware and software extensions to MONROE platform. Finally, it summarizes the hardware related issues that led to delay in the deployment of the mobile nodes and corresponding project extension by 5 months. EC accepted the amendment in March 2017.
- **Coordination of all technical and non-technical activities carried out in the project, and overall management of the consortium:** Project management and coordination have been smooth and no major issues have arisen. Cooperation between partners has been strong; this is attested by joint software development as well as authorship of scientific papers under submission. First year report has been prepared and delivered on time. 1st interim payment has been received from EC in August 2016 and distributed to the partners in September 2016. The amendments of GA have been completed in

close collaboration among the partners. Quarterly management reports were leaded by SRL and partners actively contributed with information. Finally, plenary meetings were organized and information was spread to all the partners in good time.

- **Monitoring the quality of the outputs of the project, and ensuring they are well aligned with the project's objectives and according to the work plan:** Procedures for ensuring the quality of deliverables were defined early on (Deliverable 6.1). SRL has been involved in the deliverable production process, has supported the consortium in this process, and has acted as final reviewers of all reports.

- **Ensuring communication, collaboration, consensus and information sharing between partners, and maintaining private online collaboration tools:** We organized two plenary face-to-face meeting with the help of the partners:

  1. 15-17 June, 2016: 3rd plenary meeting was held in Oslo, Norway at Simula's premises, lasting two and a half days. All WP have been discussed during the meeting. Focus on deliverables D2.2 and D4.1: status and plans.
  2. 24-28 October, 2016: 4th plenary meeting was held in Madrid, Spain at IMDEA's premises, lasting three days. All WP have been discussed during the meeting. Focus on experiments and maintenance.

  In addition to this, 1st year review meeting with EC and all the partners was organized in Oslo May 24, 2016.

  Finally, monthly teleconferences had been held since April 2015, coordinated by SRL. During the first 6 months, we held telcos once a month. Later on we have increased the frequency of the telcos and held them bi-weekly. Starting from Q8, consortium has also decided to hold maintenance hours every week in order to coordinate the platform maintenance efforts. A set of online collaboration tools was set up by the launch time of the project: a wiki, mailing lists, a repository with a version-control system and a teleconferencing system (Deliverable 6.1). SRL administers all these tools.

- **Coordinating strategic collaboration with other on-going projects and activities of relevance to MONROE:** We have initiated collaboration activities with two H2020 projects: NEAT and MAMI. Both of these projects are planning to use MONROE testbed to run their experiments. NEAT will use the MONROE platform to test relevant NEAT components, especially the multihoming feature of MONROE is crucial for NEAT. The real-network measurements produced in MONROE will provide input to NEAT system for smart interface selection in multihomed networks. MAMI project investigates middleboxes in internet and will use MONROE platform to specifically test the middlebox behaviour in mobile networks. Furthermore, we have started collaboration with FP7 FLEX project, which provides a controlled environment for LTE experiments. Our collaboration with FLEX will lead to a better understanding of mobile networks considering both controlled (FLEX) and operational (MONROE) settings.

- **Maintaining the liaison with the Advisory Board (AB):** The AB was engaged via early sharing of deliverables. The AB members have provided useful comments and feedback that the project has considered. The consortium will arrange a second plenary meeting with the AB in June, 2017.

- **Communicating to third party external users to ensure fulfillment of users:** We have announced both MONROE open calls during the first two years of the project. The first MONROE open call was announced in December 2015 and the second one in November 2016. In order to communicate with the external users, we established an email address: info@monroe-project.com. We have received many emails to this address during the open calls. Consortium members were also active disseminating the open call in different venues (please see the dissemination activities).

8 Work Package 7

WP7 Dissemination and Exploitation has the goal of establishing MONROE as the European benchmark platform for MBB measurements and testbed activities. This WP includes activities to ensure impact and future sustainability of MONROE and is organized in 3 tasks. Task 7.1 focuses on the general dissemination of MONROE features and results. Task 7.2 focuses on the specific management, dissemination and exploitation of the OPEN data resulting from the measurements, and task 7.3 focuses on sustainability. All three tasks have been active during the second year.

8.1 Progress in Task 7.1

Tools and guidelines: Consortium developed in year 1 a set of tools, templates and guidelines in order to coordinate and ensure swift dissemination and exploitation. Website (www.monroe-project.eu), including news feed (blog) and a twitter page seem to give good publicity. All other collaboration tools make us work efficiently. We use Github to share and collaborate on code (incl. open source), and we use bitbucket to collaborate on papers and reports. We share publication results and datasets as open reproducible data through zenodo.org. We have created an internal wiki page, which, among other things, serves as a tool to manage dissemination activities. In the following we present some dissemination activities from year 2 of the project.

EC events and projects: It is important for MONROE to be visible in the EC community and contribute to others' experiences. CWY and SRL participate and contribute to the regular FIRE WG teleconferences. In addition, we follow closely and collaborate with the project Mapping of Broadband Services. CWY represented MONROE at the 2nd Stakeholder's Workshop: Mapping of Broadband Services (Policy makers, industry) in December 2016. A follow up from this event will be meetings with some European regulators that have plans for similar platforms.

KAU and CWY are also collaborating closely with the H2020 NEAT project, and they will provide support for the NEAT library in MONROE nodes. This enables all users of MONROE to use the NEAT transport features as an alternative to the default features. SRL is collaborating closely with H2020 MAMI project where protocol and middlebox measurements are run and analyzed for MBB networks using MONROE.

Industrial and societal outreach: Industrial and societal outreach is key to place MONROE as a benchmark platform in the continent. In year 2, we did the following presentations.

- Presentation of MONROE status for the Ministry of Transport and Communication in Norway by SRL and CWY. Oslo, March 2016.
- CWY held an industrial presentation: Keynote and panel debate on 1st Nordic conference on ICT - 5G impact at Telenor headquarter. CWY presented MONROE objectives and plans for executives from operators, vendors and service providers. Oslo, April 2016.
- KAU held an industrial presentation: Compare (local industry) Lunch Seminar. April 2016.
- KAU presented MONROE for READY project plenary. April 2016.
- IMDEA presented MONROE platform for the members of the NOTRE project (Network for Social Computing Research). June 2016.

- IMDEA presented MONROE platform to Accenture Innovation. August 2016.
- IMDEA presented MONROE platform to the Regional Government of Madrid and Roche. November 2016.
- SRL held a Platform presentation at e-Wine Workshop. November 2016.
- IMDEA held a presentation to the Secretario de Estado of Information Society and Digital Agenda from the Spanish Ministry of Energy, Tourism and Digital Agenda. January 2017.
- SRL had a meeting with the Norwegian Regulator. January 2017.

Media:

For promoting the open calls, partners' web pages and twitter have been used to disseminate practical information. The project has also produced some business cards with open call info that have been handed out on workshops, meetings and conferences. Both MONROE open calls have been widely disseminated through MONROE project website:

- https://www.monroe-project.eu/press-release-monroe-announces-first-open-call-with-a-total-budget-of-eur-18-million-for-external-experimenters/ - 2016
- https://www.monroe-project.eu/first-monroe-open-call-application-statistics/ - 2016
- https://www.monroe-project.eu/first-monroe-open-call-results/ - 2016
- https://www.monroe-project.eu/first-user-workshop-and-consortium-meeting-held-in-oslo/ - 2016
- https://www.monroe-project.eu/second-call-for-experimenters-is-open-for-applications/ - 2016
- https://www.monroe-project.eu/monroe-open-call-2-statistics/ - 2016
- https://www.monroe-project.eu/second-monroe-open-call-results/ - 2017

Press release about the open calls is published on the website of some partners (KAU):

- https://www.kau.se/index.php/en/cs/news/great-interest-monroe-first-open-call-external-users - March 2016
- https://www.kau.se/en/cs/news/new-experimenters-accepted-monroe-measurement-platform - February 2017

The information about MONROE have been picked up by other web pages:

- Released on SRS (associated partner from 1st open call) website:http://www.softwareradiosystems.com/news_2016_08_08_monroe/ -August 2016
- Announced on Twitter: https://twitter.com/H2020MONROE/status/829343528388063232 - February 2017

Scientific community outreach: MONROE is a Research and Innovation action. Therefore, contributions to scientific community are a big component of the project. To this end, the consortium has produced papers describing the platform and has received a best paper award for one of them.

IMDEA Networks MONROE project research team has won the Best Demo Award at WoWMoM 2016, the 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks, which was held from June 21st to the 24, in Coimbra, Portugal. WoWMoM is one of the main international venues for the presentation and exchange of ideas about the top networking research challenges: http://www.networks.imdea.org/whats-new/news/2016/wowmom2016-best-demo-award-goes-imdea-networks-monroe-project-research-team - July 2016

Published papers and presentations:

- MONROE: Measuring Mobile Broadband Networks in Europe. Chapter accepted to be published in FIRE Book. (ALL, August 2016)
- SRL Paper: M. Fida, A. Lutu, M. Marina, and O. Alay. "ZipWeave: Towards Efficient and Reliable Measurement based Mobile Coverage Maps." November 2016. Accepted for publication in IEEE INFOCOM 2017.
- SRL Paper: Profiling Mobile Broadband Coverage. Author list: Andra Lutu, Yuba Raj Siwakoti, Ozgu Alay, Dziugas Baltrunas and Ahmed Elmokashfi. Traffic Monitoring and Analysis (TMA). Belgium, April 2016.
- SRL Paper: The good, the bad and the implications of profiling Mobile Broadband Coverage. Author list: Andra Lutu, Yuba Raj Siwakoti, Ozgu Alay, Dziugas Baltrunas and Ahmed Elmokashfi. Elsevier Computer Networks. October 2016.
- SRL Paper: Investigating Packet Loss in Mobile Broadband Networks under Mobility. Author list: D. Baltrunas, A. Elmokashfi, A. Kvalbein and O. Alay. IFIP Networking. Austria, May 2016.
- Poster by POLITO: Video Upload from Public Transport Vehicles using Multihomed Systems. Author list: Ali Safari Khatouni, Marco Ajmone, Marco Mellia. IEEE INFOCOM Student Workshop. San Francisco, April 2016.
- Paper by POLITO: Delay Tolerant Video Upload from Public Vehicles. Author list: Ali Safari Khatouni, Marco Ajmone, Marco Mellia. SmartCity 2016, The 2nd IEEE INFOCOM Workshop. San Francisco, April 2016.
- Paper by IMDEA: OWL: a Reliable Online Watcher for LTE Control Channel Measurements. Author list: Nicola Bui, Joerg Widmer. All Things Cellular (ATC'16) in Conjuction with ACM Mobicom at City University New York . October, 2016.
- Conference presentation and platform demo by IMDEA at IEEE WoWMoM. June 2016. **Best Demo Award.** (paper by ALL, demo by IMDEA).
- Conference presentation and platform demo by KAU at ACM WiNTECH'16 in Conjuction with ACM Mobicom at City University New York. October 2016. (paper by ALL, demo by KAU).
- Conference presentation/demo by KAU at workshop on All Things Cellular (ATC'16) in Conjuction with ACM Mobicom at City University New York. October, 2016. (paper by ALL, demo by KAU).

- KAU presented MONROE demo for KAU computer science researchers and international advisory board. September 2016.
- KAU held a presentation and presented platform for KAU CS research seminar/colloqium. September 2016.

Work in Progress:

- MONROE Experience paper (ALL): "Experience: An Open Platform for Experimentation with Commercial Mobile Broadband Networks", submitted to ACM MobiCom 2017, the Annual International Conference on Mobile Computing and Networking, March 2017 (currently under review, notification of acceptance due June 8, 2017).
- MONROE Dataset paper (ALL): "A large-scale data collection framework for commercial mobile broadband networks: The MONROE case", submitted to ACM SIGCOMM 2017 Workshop on Big Data Analytics and Machine Learning for Data Communication Networks (Big-DAMA 2017), MArch 2017 (currently under review, notification of acceptance)
- MONROE Experiments: A. Safari Khatouni, M. Mellia, M. Ajmone Marsan, S. Alfredson, J. Karlsson, A. Brunstrom, O. Alay, A. Lutu, C. Midoglu, V. Mancuso. "Speedtest-like Measurements in 3G/4G Networks: the MONROE Experience." Submitted to the 29th International Teletraffic Conference (currently under review, notification of acceptance due May 14, 2017)
- MONROE Experiments (together with the MAMI project): I. Learmont, A. Lutu, G. Fairhurst, D. Ros, O. Alay. "Path Transparency Measurements from the Mobile Edge with PATHspider." Submitted to IEEE/IFIP Workshop on Mobile Network Measurement (MNM'17), held in conjunction with TMA'17 (currently under review, notification of acceptance due May 5, 2017)
- MONROE System paper (ALL): "MONROE: A Large-scale Platform for Mobile Broadband Measurements and Experiments", submitted to IEEE Communication Magazine, Network Testing and Analytics Series, August 2016 (rejected, working for a new submission to the same venue).

External Users Papers:

- Maril-in-Monroe paper: E. Atxutegi, J. O. Fajardo, E. Ibarrola and F. Liberal, "Assessing Internet performance over mobile networks: From theory to practice," 2016 ITU Kaleidoscope: ICTs for a Sustainable World (ITU WT), Bangkok, 2016, pp. 1-8. doi: 10.1109/ITU-WT.2016.7805705
- Maril-in-Monroe paper: Eneko Atxutegi, Andoni Izurza, Fidel Liberal, Ake Arvidsson, Karl-Johan Grinnemo, Anna Brunstrom. "Open issues in the interaction between TCP and the current and future mobile networks." Submitted to IEEE/IFIP Workshop on Mobile Network Measurement (MNM'17), held in conjunction with TMA'17 (currently under review, notification of acceptance due May 5, 2017)
- Mobi-QoE paper: Michael Seufert, Nikolas Wehner, Florian Wamser, Pedro Casas, Alessandro D'Alconzo, Phuoc Tran-Gia. "Unsupervised QoE Field Study for Mobile YouTube Video Streaming with YoMoApp." Submitted to the 9th International Conference on Quality of Multimedia Experience (QoMEX 2017) (currently under review, notification of acceptance due April 25, 2017)
- Mobi-QoE paper: Michael Seufert, Ondrej Zach, Martin Slanina, Phuoc Tran-Gia. "Unperturbed Video Streaming QoE Under Web Page Related Context Factors." Submitted to the 9th International Conference on Quality of Multimedia Experience (QoMEX 2017) (currently under review, notification of acceptance due April 25, 2017)

- Mobi-QoE paper: Anika Schwind, Michael Seufert, Ozgu Alay, Pedro Casas, Phuoc Tran-Gia, Florian Wamser. "Concept and Implementation of Video QoE Measurements in a Mobile Broadband Testbed." Submitted to IEEE/IFIP Workshop on Mobile Network Measurement (MNM'17), held in conjunction with TMA'17 (currently under review, notification of acceptance due May 5, 2017)
- SOMETIME paper: Giuseppe Aceto, Valerio Persico, Antonio Pescape, Giorgio Ventre. "SOMETIME: SOftware defined network-based Available Bandwidth MEasuremenT In MONROE." Submitted to IEEE/IFIP Workshop on Mobile Network Measurement (MNM'17), held in conjunction with TMA'17 (currently under review, notification of acceptance due May 5, 2017)
- MONROE-SOPHIA paper: Paul Sutton and Ismael Gomez. "MONROE-SOPHIA - A Software Radio Platform for Mobile Network Measurement." Submitted to IEEE/IFIP Workshop on Mobile Network Measurement (MNM?17), held in conjunction with TMA'17 (currently under review, notification of acceptance due May 5, 2017)
- Monroe-LTE paper: Pedro M. B. Torres, Paulo Marques, Hugo Marques, Tiago Alves, Jorge Ribeiro, Rogério Dionisio, "Data analytics for forecasting cell congestion on LTE networks." Submitted to IEEE/IFIP Workshop on Mobile Network Measurement (MNM'17), held in conjunction with TMA'17 (currently under review, notification of acceptance due May 5, 2017)
- NAPPLYTICS paper: Irene Alepuz, Jorge Cabrejas, Jose F. Monserrat, Alvaro G. Perez, Gonzalo Pajares, Roberto Gimenez. "Use of Mobile Network Analytics for Application Performance Design." Submitted to IEEE/IFIP Workshop on Mobile Network Measurement (MNM'17), held in conjunction with TMA'17 (currently under review, notification of acceptance due May 5, 2017)
- PREC paper: Ana Custura, Andre Venne, Gorry Fairhurst. "Exploring DSCP modification pathologies in mobile edge networks." Submitted to IEEE/IFIP Workshop on Mobile Network Measurement (MNM'17), held in conjunction with TMA'17 (currently under review, notification of acceptance due May 5, 2017)

8.2 Progress in Task 7.2

The goal of this task is to ensure that the dataset generated by MONROE is accessible and useful for the maximum number of interested researchers. In order to accomplish this, we have agreed on three parallel paths:

1. A secondary copy of the MONROE (Cassandra) database has been prepared by IMDEA. This database is accessible for all members of the first open call users. However, to improve the safety and usability of the machine, access credentials are granted on demand.
2. Every day a CSV dump of all the tables is generated and stored in a place accessible through the same web server that provides access to the user interface and scheduler API. Hence, a valid Fed4Fire user certificate is needed to access the dump files. This is a temporary measure that will probably be waived during the next periods, when our dataset is fully mature.
3. The consortium is evaluating other long-term solutions to store our dataset for posterity, considering the solutions offered by Zenodo as the most likely candidate.

An interim report on the dissemination activities has been provided as part of Deliverable 7.1, which was due month 18. This report outlined the plans and current status for dissemination activities, guidelines and prioritizations with milestones. Milestone 3 was achieved by the completion of dissemination plan and the launch of the project website during the first year. Milestone 12 Open Data was due in month 24 and is achieved according the plans. Initial Open Data is available in project website and public repositories.

8.3 Progress in Task 7.3

Sustainability of the platform is of high importance for all partner institutions. In the plenary meetings, we have had sustainability planning as an important topic. We are working on competitive analysis, value propositions and business models that will result in brochure and be presented and discussed in a stakeholders workshop. The goal was to have a first version ready in March 2017, with a follow-up workshop in June 2017. However, due to unexpected HW challenges, we have postponed the conclusion of these work to August 2017 with a workshop in November 2017.

Recently, we finalized an analysis of alternative initiatives and products in order to identify our competitive advantages as background for value propositions. Out of the 23 initiatives we analyzed, 16 where companies, 4 institutions and 3 regulators. 17 of them targeted operators, vendors or enterprises, 6 targeted consumers while 4 targeted research institutions. 7 were based on crowdsourcing apps while 5 used specialized HW equipment.

The list of competitors presented a wide variety of actors. Some crowdsource apps focused on capturing speed and some metadata to map coverage. Others again did QoE/QoS measurements for specific applications/services. The specialized HW focused on measuring operator specific configurations with limited measurement points. MONROE is more generalized, affordable and spread out than the HW and operator specific tools. MONROE is dedicated and not shared with other data as the crowdsourcing solution, and MONROE is experimenter-initiated and not user initiated as most of the smartphone apps.

The analysis shows that MONROE has some valuable competitive advantages that will form the basis for value propositions and business models.

- MONROE can run any experiment that can run in Linux. They can be develeoped and deployed in a day if a must. No other platforms can offer this.
- It is a dedicated long-term platform that enables measurement of long-term evolution of operators/paths from same location.
- The possibility to test with and measure three MBB operators (+ WiFi and Ethernet) from same location over time is unique.
- MONROE can provide more detailed metadata than most crowdsourced apps.
- The host networks that provide us with nodes spread out in Europe, on trains, busses and trucks seems to be unique.

Another source for sustainability is new EU projects. Partners have engaged in applying new projects that might help funding the platform. Most of them were with 5G focus. We have also been contacted by institutions requesting to buy nodes and licensing the platform software. When we complete our work on business models, we will promote this widely.

We have also highlighted sustainability as an important evaluation criteria in the **open calls**: "For the second call, proposals that have a solid plan to contribute to the sustainability of the MONROE platform will be prioritized among equally scored proposals."

There where two projects from the first open call that had clear plans for business and sustainability. One selling data and knowledge based on a combination of crowd sourced data and MONROE, and the other target operators with a RIPE like model. From the second open call five funded projects presented relevant plans for business and sustainability. Some suggest to use MONROE as an additional product/platform in business they already have nationally. One suggests a Sam knows model for more detailed ISP analysis like NAT effects. Another finds the data valuable for selling to operators, while another will license to research institutions and SMEs. We will keep a close dialogue with all of them on how to succeed with their ideas.

9 Project Impact

The expected impact and measures presented in the DoA are still relevant for the work and plans in the project. However, it is worth mentioning that the proposals for open calls exceeded our expectations in terms of the variety of the applications that MONROE can support. Especially the accepted proposals target to solve different problems in different verticals spanning form energy domain (smartgrid) to vehicular communications to health domain. Furthermore, the projects in the communication and networking domain address very important problems that are faced by operators in their daily operations and provide unique and innovative solutions to the operators for them to improve their networks. We are working closely with our external users to further increase the project's and platform's impact in different verticals.

Although we reduced the total number of nodes due to the hardware related issues, we believe the extensions provided by the external users compensate for this by not only increasing the number of nodes but also providing diversity. Moreover, the delays in the large scale deployment will be compensated by the 5 months extension to the project. Therefore, we believe the project can fullfill its targeted impact for large-scale monitoring of MBB networks.

10 Update on the Exploitation and Dissemination Plan (if applicable)

Deliverable 7.1, presented and updated dissemination plan. We presented a goal of having stakeholder brochures ready in March in order to plan for a workshop in June. Due to unforeseen challenges with hardware and late deployment, we will launch brochures in August and target a stakeholders workshop in November. It is important to generate considerable amount of data and knowledge before completing these activities.

We are planning to contribute to the standardization bodies such as IETF. We have initial presentations in different workgroups of IETF and during the last year of the project, we will present the results of the project, especially in the IETF Large-scale Measurements for Broadband Performance (LMAP) Working Group, IETF MPTCP Working Group, and IRTF ICCRG Research Group.

11 Update of the Data Management Plan (if applicable)

No updates have been made to data management since D6.2. MONROE will use Zenodo for sharing, archiving and identifying (via DOIs) the open data. Github is the code repository chosen for all open-source software released by the project. The implementation of MONROE system has already been shared as open source in github. Scientific publications will be shared on Zenodo together with results files as well as the scripts that are used to produce these results in order to ensure reproducibility.

12 Follow-up of Recommendations and Comments from Previous Review(s) (if applicable)

In the first year review of the project, we have received some suggestions from the review panel. Below, we list the main comments and how we addressed them.

1. **Platform Description and Experience:** The reviewers have clearly mentioned that the project needs to focus on putting together an effective and stable system first and then write a paper summarizing our experience while building the platform. The consortium worked really hard to eliminate many instabilities in the platform. We have decided to upgrade the node hardware to eliminate majority of the instabilities. All our experience is summarized in a double blind ACM Mobicom paper submission that can be found in Appendix II. This submission is currently under review and considers different aspects of the platform and provides the take aways which we believe is very valuable to the community.

2. **Use Cases:** The reviewers have mentioned that the project needs to sort out the use cases of the platform and what facilities like ours can achieve in the future. To address different use-case, we are working on different paper submissions. Two notable submissions in this direction are: (i) ACM Sigcomm BigDama paper that focuses on how the MONROE dataset can enable different use-cases and (ii) ITC paper that focuses on whether speedtest-like measurements are easy to conduct in mobile networks and how a platform such as MONROE enables dissecting the mobile networks including detection of proxies in these networks. We are currently working on another use case, web application performance, and we target to submit this work to ACM IMC 2017 or ACM Conext 2017.

3. **Tstat integration:** The reviewers mentioned that the Tstat has to be integrated to the project in a much better way. Throughout the second year, we have worked on this integration. The current implementation provides Tstat as metadata to the users of the platform and all the statistics collected from Tstat is provided in MONROE database.

4. **Exploration of new sources of funding:** The reviewers expressed that the project should prioritize the sustainability of the platform, especially how to get money from the private sector. Please see Section 8.3 for the details of the work carried out in sustainability related activities.

5. **Building a community:** The reviewers suggested us to work closely with external users and build a community around the platform. Through the open calls, the project already supports 27 external users. Each project has a patron closely following them. We also have a mailing list where all the users can interact and get help from the consortium but also from the other users. When the platform is open, we will reach to the users that did not get funding. We are also working with other EU projects such as MAMI to arrange workshop on Mobile Network Measurement at TMA'17, reaching out a wider community. Furthermore, despite the hardware related problems, the project has been productive in terms of papers, including the papers of the external users (listed in Section 8.2).

13 Deviations from Appendix I (if applicable)

13.1 Tasks

Towards the end of the second year, due to hardware and software issues, the design of the MONROE platform has undergone a significant change. The issues summarized in Section 3 suggested the strong need to modify the platform design, particularly to avoid clear safety risks. The tasks impacted from these changes are summarized below:

Task 1.2 and Task 1.3: With the submission of D1.3, these tasks have been concluded. We foresee small updates to the MONROE SW due to updates to the nodes. All the necessary updates to the *User Access and Experiment Scheduling System* and *Data Collection, Storage and Visualization System* will be carried out under the Maintenance work package, WP4.

Task 2.2: Due to the above mentioned reasons, the deployment of the nodes are ongoing. The main focus of this task is the transition from the old hardware system to the new dual-APU system. We prioritize the mobile nodes and upgrade the stationary nodes that are operational in time. The deployment is expected to be concluded at the end of June, 2017. To accommodate the delays in the deployment of the nodes, the project period has been extended by 5 months.

13.2 Use of resources

Regarding the use of resources, there are four deviations with respect to the DoA:

1. SRL has overspent resource effort in WP1 with respect to the original plan in the first two years (resulted in 5.5 PM overspending). This overspending was partially due to problems we experienced in the prototype implementation and additional time spent to identify and debug the problems in the initial node design. Moreover, during the software development and testing, as well as the experiment design, PhD students and young engineers help our senior engineers and researchers. The effort associated to those lower-experienced personnel was necessarily higher than the one required by senior engineers and researchers. However, in terms of overall cost of personnel, SRL's overspending in the first two years is much less accentuated than PM overspending.

2. IMDEA has used 5.5 PM more than planned over the first 2 years. However, in the second year IMDEA has used less than what had been planned exclusively for Y2. The reason is that IMDEA has carried out bulk work in Y1 and we have anticipated some activities from Y2 to Y1. In addition, this required the use of more engineers and the involvement of some students to speedup the design (and optimize the management) of Database replicas and user access portal.

3. Due to the problems with the node hardware and the delay in the deployment of the nodes, KAU has delayed much of the planned experimental work to be able to perform the measurements with a stable platform. These resources will instead be used during the third year of the project.

4. POLITO has claimed additional 7 PM for undergraduate researcher as an adjustment for Y1. It was based on the recent changes to the Model Grant Agreement introduced on 27th February 2017. The difference was added to the overall cumulative values for Y1 and Y2 in Figure 1.

5. Majority of Telenor's effort is in WP3 Experiments. Due to the hardware related problems, Telenor has not spent the efforts as planned and will re-allocate this effort in the third year to run extensive measurements campaigns and to analyze the results.

13.2.1 Unforeseen subcontracting (if applicable)

Not applicable.

13.2.2 Unforeseen use of in kind contribution from third party against payment or free of charges (if applicable)

Regarding open calls, there is no budget for external reviewer's fee and, therefore, external reviewers costs will represent an in-kind contribution free of charge to MONROE.

Table 4: Actual versus planned efforts during Year 2.

Participant	WP1		WP2		WP3		WP4		WP5		WP6		WP7		Total	
	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual
SRL	9	13,90	1,82	0,97	5,5	5,67	3,86	1,89	2,66	0,58	6	3,96	1,34	0,94	30,18	**27,91**
IMDEA	13	5,1	0,91	0,6	6,5	5,65	2,7	2,61	4	2,75	0	0	1,8	1,71	28,91	**18,42**
KAU	5,5	3,75	0,5	0,53	5,5	0,15	1,29	0,29	4,45	1,83	0	0	1	0,25	18,24	**6,80**
POLITO	6	4,22	0,91	1,55	4,50	4,93	0,86	1,37	0,89	1,21	0	0	1,00	1,59	14,16	**14,87**
NXW	2,5	3,09	13,00	15,91	0	0,00	2,14	1,87	0	0	0	0	1,34	0,67	18,98	**21,54**
Celerway	2,5	2,03	1,82	0,6	1	0,3900	0	0	0	0	0	0	1,37	1,17	6,69	**4,1900**
Telenor	1	0,04	0	0	4,5	3,46	0	0	0	0	0	0	1	0,33	6,5	**3,83**
Total	**39,50**	**32,13**	**18,96**	**20,16**	**27,50**	**20,25**	**10,85**	**8,03**	**12,00**	**6,37**	**6,00**	**3,96**	**8,85**	**6,66**	**123,66**	**97,56**

Table 5: Cumulative values of actual versus planned efforts, since the project started.

Participant	WP1		WP2		WP3		WP4		WP5		WP6		WP7		Total	
	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual	Planned	Actual
SRL	18	27,66	4	2,41	5,5	5,67	3,86	1,89	3,33	0,79	12	10,57	2,67	2,36	49,36	**51,35**
IMDEA	26	31,64	2	0,78	6,5	5,65	2,7	2,61	4	2,75	0	0	3	4,7	42,7	**48,13**
KAU	11	10,92	1	1,34	5,5	0,26	1,29	0,29	5,56	2,14	0	0	2	0,61	25,35	**15,56**
POLITO	12	12,51	2,00	2,68	4,50	4,93	0,86	1,37	1,11	1,42	0	0	2,00	2,54	22,47	**25,45**
NXW	5	5	26,00	26,43	0	0,00	2,14	1,87	0	0	0	0	2,67	0,96	35,81	**34,26**
Celerway	5	5,28	4,00	4,11	1	0,3900	0	0	0	0	0	0	2,70	3,07	12,7	**12,85**
Telenor	2	1,63	0	0	4,5	3,46	0	0	0	0	0	0	2	0,67	8,5	**5,76**
Total	**79,00**	**94,64**	**39,00**	**37,75**	**27,50**	**20,36**	**10,85**	**8,03**	**14,00**	**7,10**	**12,00**	**10,57**	**17,04**	**14,91**	**196,89**	**193,36**

Y2 PM - PP.pdf Y2 PM - PP.pdf Y2 PM - PP.pdf

Figure 1: Use of resources tables.

Appendix I

Status of the Deliverables and Milestones

During the first two years, SRL ensured (a) the verification of project deliverables by the appropriate persons in the project management structure, and (b) the timely submission of such deliverables to the EC Scientific Officer. To this end, 13 deliverables have been submitted and 12 milestones have been achieved as tabulated in the below Tables.

Furthermore, Quarterly Management Reports are generated and submitted to the Commission, containing basic facts and figures on a per-quarter basis, including brief updates on work progress and use of human resources. Since the beginning of the project, 8 QMRs have been submitted to summarize the status of the project.

Number	Deliverable title	Due date	Delivery date
D6.1	Report on project management tools	June 1st, 2015	June 1st, 2015
D6.2	Data Management Plan	September 1st, 2015	September 1st, 2015
D1.1	Report on use cases	September 1st, 2015	September 14th, 2015
D1.2	System Design and Prototype implementation	March 1st, 2016	March 1st, 2016
D2.1	Selection and design of the new node	March 1st, 2016	March 1st, 2016
D6.3	First year report	March 1st, 2016	with periodic report in April
D2.2	Node Deployment	June 1st, 2016	September (agreed with PO)
D4.1	Maintenance routines	September 1st, 2016	September 15th(agreed with PO)
D7.1	Interim dissemination report	September 1st, 2016	September 15th(agreed with PO)
D2.3	Robust node recovery method	December 1st, 2016	December 1th, 2016
D1.3	Final implementation of software	March 1st, 2017	February 28th, 2017
D3.1	Experimental SW as EaaS	March 1st, 2017	February 28th, 2017
D5.1	User manual and initial user's experiences	March 1st, 2017	February 28th, 2017

Number	Milestone	Due date	Achieved date
MS1	System Design and Proof of Concept	September 1st, 2015	September 1st, 2015
MS2	HW selection	September 1st, 2015	September 1st, 2015
MS3	Dissemination Plan and public website	September 1st, 2015	September 1st, 2015
MS4	Prototype implementation	March 1st, 2016	March 1st, 2016
MS6	Experimenters in phase I ready to use the platform	June 1st, 2016	June 1st, 2016
MS7	Maintenance Routines	September 1st, 2016	September 1st, 2016
MS8	Phase I experimenters feedback to WP1	September 1st, 2016	September 1st, 2016
MS9	Final implementation	March 1st, 2017	March 1st, 2017
MS10	Experimental SW and Tools for External Users	March 1st, 2017	March 1st, 2017
MS11	Experimenters in phase II ready to use the platform	March 1st, 2017	March 1st, 2017
MS12	Open Data	March 1st, 2017	March 1st, 2017

Appendix II

We have gathered our experience with the MONROE platform in a paper currently under submission at Mobicom 2017, for the "Experience" track. We attach to this document the paper titled "Experience: An Open Platform for Experimentation with Commercial Mobile Broadband Networks", which we submitted on March 16, 2017. The paper complies with the double-blind requirements of the conference (i.e., "OpenMBB" is the anonymized name for the MONROE platform).

Experience: An Open Platform for Experimentation with Commercial Mobile Broadband Networks

ABSTRACT

Open experimentation with operational Mobile Broadband (MBB) networks in the wild is currently a fundamental requirement of the research community in its endeavor to address the need of innovative solutions for mobile communications. Even more, there is a strong need for objective data about stability and performance of MBB (e.g., 3G/4G) networks, and for tools that rigorously and scientifically assess their status. In this paper, we introduce the OpenMBB measurement platform: an open access and flexible hardware-based platform for measurements and custom experimentation on operational MBB networks. The OpenMBB platform enables accurate, realistic and meaningful assessment of the performance and reliability of 11 MBB networks in Europe. We report on our experience designing, implementing and testing the solution we propose for the platform. We detail the challenges we overcame while building and testing the OpenMBB testbed and argue our design and implementation choices accordingly. We describe and exemplify the capabilities of the platform and the wide variety of experiments that the system already supports.

1. INTRODUCTION

Mobile broadband (MBB) networks have become the key infrastructure for people to stay connected everywhere they go and while on the move. Society's increased reliance on MBB networks motivates researchers and engineers to enhance the capabilities of mobile networks by designing new technologies to cater for plethora of new applications and services, growth in traffic volume and a wide variety of user devices. In this dynamic ecosystem, there is a strong need for both open objective data about the performance and reliability of different MBB operators, as well as open platforms for experimentation with operational MBB providers.

In this paper, we introduce ***OpenMBB****: the first open access hardware-based platform for independent, multihomed, large-scale experimentation in MBB heterogeneous environments.* OpenMBB comprises a large set of custom hardware devices, both mobile (e.g., via hardware operating aboard public transport vehicles) and stationary (e.g., volunteers hosting the equipment in their homes), which are multihomed to three different MBB operators using commercial grade subscriptions[1].

Evaluation of network performance, assessing the quality experienced by end users and experimenting with novel protocols require thorough systematic repeatable end-to-end measurements. While existing experimental platforms, such as Planetlab [18], RIPE Atlas [20] and CAIDA Ark [3], meet these requirements, they are limited to fixed broadband networks and are not multihomed. OpenMBB is an one-of-a-kind platform that enables controlled experimentation with different commercial MBB providers. It enables users to run custom experiments and to schedule experimental campaigns to collect data from operational MBB and WiFi networks, together with full context information (metadata). For example, OpenMBB can accommodate performance evaluation of different applications (e.g., web and video) over different networks or testing different protocols and solutions under the same conditions.

Objective performance data is essential for regulators to ensure transparency and the general quality level of the basic Internet access service [14], especially in light of an evolution of service offerings beyond the best-effort traffic mode, including a balanced approach to net neutrality. Several regulators responded to this need with ongoing nationwide efforts [6]. Often, they do not open the solutions to the research community to allow for custom experimentation, nor do they grant free access to the measurement results and methodology. OpenMBB aims to fill this gap and offers free access to custom experimentation with commercial MBB networks.

A common alternative to using controlled testbeds such as OpenMBB is to rely on end users and their devices to run tests by visiting a website [16] or running a special application [12]. The main advantage of such crowdsourcing techniques is scalability: it can collect millions of measurements from different regions, networks and user equipment types [9]. However, repeatability is challenging and one can only collect mea-

[1] At the time of writing, we are in the process of deploying the devices to their final hosts. We are scheduled to complete this phase by June 2017.

surements at users' own will, with no possibility of either monitoring or controlling the measurement process. Mostly due to privacy reasons, crowdsourced measurements do not always provide important context information (e.g., location, type of user equipment, type of subscription, and connection status (2G/3G/4G and WiFi)). Finally, with this approach it is challenging to test novel applications and services since innovative solutions might require configuration changes (e.g., customized kernels). OpenMBB is complementary to crowdsourcing approaches and the control over the measurement environment tackles the shortcomings of crowd data. Furthermore, OpenMBB supports the deployment of different applications and protocols, and enables benchmarking tools and methodologies.

In the rest of the paper, we report on our experience designing, implementing and using the platform. We detail the design considerations and demonstrate the versatility of our approach (Section 2). We explain how we cater for the requirements of experimenters and enable them to deploy myriad measurements on operational commercial MBB networks. The OpenMBB measurement node (hereinafter, the node or the OpenMBB node) sits in the center of the system and is the most important element, conditioning the proper functionality of the platform and the associated measurement system. Next, we describe our experience with the OpenMBB system implementation and detail the hardware selection for the OpenMBB measurement node (Section 3). We forged the node to be flexible and powerful enough to run a wide range of measurement and experimental tasks, including demanding applications like adaptive video streaming. In the same time, we ensured that the node software design translates into a robust implementation (Section 4) that is also easily evolved and upgraded in order to sustain the most recent technological innovations. We further present the user access and scheduling solution we offer experimenters for exploiting the available resources of the platform in a fair manner (Section 5). Finally, we demonstrate that the OpenMBB system is a fitting solution to conduct a wide range of experiments over commercial cellular networks. To showcase its capabilities, we describe different categories of experiments OpenMBB supports (Section 6).

2. SYSTEM DESIGN

Throughout the design process of the OpenMBB platform, we interacted with the users of the platform (e.g., universities, research centers, industry and SMEs) and collected their feedback on requirements for the platform functionality. This allowed us to gauge experimenters' expectations from such a system and based on this, we sketch the platform specifications.

2.1 Requirements

We summarize the main requirements as follows.

Large scale and Diversity: To give a representative view of the characteristics of an entire network, we need to collect measurements from a large number of vantage points. Furthermore, we should strive to collect measurements under diverse geographical settings, from major cities to remote islands.

Mobility: Mobility is what makes MBB networks unique compared to other wireless networks. To provide insight into the mobility dimension of MBB networks, it is imperative that the platform integrates a deployment under realistic mobility scenarios.

Fully programmable nodes: To accommodate the wide range of experiments users contemplate to run on the platform, we should forge measurement devices that are flexible, powerful and robust.

Multihoming support: To compare different mobile operators and/or different wireless technologies under the same conditions, the same node should connect to multiple providers in the same time (multihoming support). This further makes the platform particularly well suited for experimentation with methods that exploit aggregation of multiple connections.

Rich context information: While analyzing the measurements, context information is crucial. The platform should monitor the network conditions, the time and location of the experiment, as well as the metadata from the modems, including, for example, cell ID, signal strength and connection mode.

Easy to use platform: It is crucial to make it easy for users to access the system and deploy experiments on all or a selected subset of nodes. This requires a user friendly interface together with a well managed and fair scheduling system.

2.2 Design Overview

We shaped the main building blocks of the OpenMBB platform such that we can meet the above-mentioned requirements. Note that while implementing different components of the platform, operational aspects also impacted the design choices, which we will discuss in detail Sections 4-5. Next, we give an overview of the purpose and functionality of the main building blocks of the OpenMBB system, which we illustrate in Figure 1.

OpenMBB Node: OpenMBB operates 250 nodes[2] in 4 countries in Europe (Spain, Italy, Sweden and Norway). The measurement node is in the core of our platform and its design comprises two main notions, namely the hardware configuration, and the software ecosystem. In terms of hardware, each node has a main board that is a small programmable computer and supports (at least) 4 interfaces: three 3G/4G modems and one Wifi modem. To cover a diverse set of mobility scenar-

[2]At the time of writing, we have deployed 50 out of the 250 planned nodes. Node assembly and installation is ongoing.

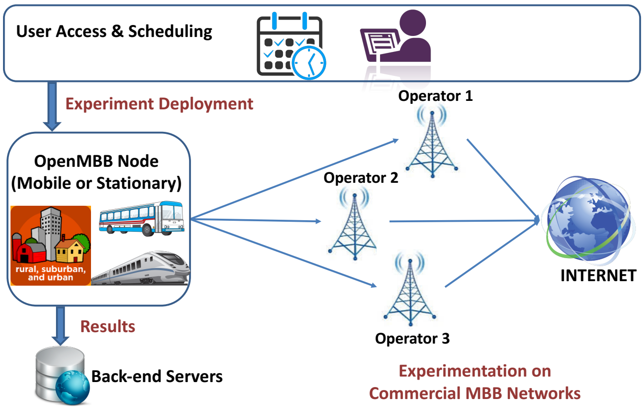

Figure 1: Main blocks of the OpenMBB platform: The OpenMBB Nodes operate in trains, buses or inside homes and connect to three commercial MBB operators in each country with OpenMBB presence. The users access the resources and deploy their experiments via the User Access and Scheduling. We periodically synchronize the measurement results to external repositories operating in the back-end.

ios, we customize a portion of the nodes (i.e., 100 out of 250 total nodes) to operate on public transport vehicles (buses and trains) and also in delivery trucks. In Section 3, we detail the choices for the node hardware implementation, our experience with running two node prototypes and our current solution.

The node software is based on a Linux Debian *Stretch* distribution to ensure compatibility with multiple hardware configurations and to enable a large set of experiments. Furthermore, especially considering the experimentation on protocols, Linux is the only operating system with sufficient hardware support for research and implementation of transport protocols due to the accessibility of the source code, flexibility and community maintenance to ensure operability with other systems. On top of the operating system, the nodes run: (i) *the management software* that performs the normal jobs expected on any mobile device, (ii) *the maintenance software* that monitors the operational status of the nodes and diminishes manual maintenance intervention and (iii) *the experimentation enablers*, that enable experiment deployment (via the scheduler client) and feed rich context information to the experiments. To provide agile reconfiguration and access for the experimenter to different software components, the experiments run in the Docker light-weight virtualized environment. This also ensures the containment of external actions in the node system. We periodically transfer the results of the experiments from the nodes to a remote repository. We further detail in Section 4 the node software ecosystem and present our evaluation of potential node internal performance overheads.

User access and scheduling: OpenMBB enables User Access to the platform resources through a user-friendly web portal that allows authenticated users to use the OpenMBB scheduler to deploy their experiments. The OpenMBB Scheduler facilitates exclusive access to the nodes (i.e., no two experiments run on the node at the same time) while ensuring fairness among users by accounting data quotas. We provide the details and the implementation choices for the user access and scheduling policies in Section 5.

3. HARDWARE IMPLEMENTATION

Given the requirements we draw from the wide range of use-cases experimenters envision for the OpenMBB platform, the measurement device needs to be small, affordable, robust, sufficiently powerful and should support the mainline Linux kernel. The size and price constraints limited us to evaluate different Single Board Computers (SBCs). There is a large amount of different SBCs available to the consumer public, with different CPU architectures and hardware configurations. However, most contain hardware requiring the use of proprietary drivers, thus restricting us to old kernels or making it impossible to compile custom kernels. We evaluated several options, including popular ones such as Raspberry Pi [19], Odroid [15], Beaglebone [1] and we selected PcEngines APU [17]. APUs integrate a 1Ghz 64 bit quad core processor, 4GB of RAM and a 16GB HDD. This provides sufficient power, storage and memory for the foreseeable future. APUs have 3 miniPCI express slots, two of which support 3G/4G modems.

Modem Selection: To multihome to 3 mobile operators and a WiFi hotspot, we equipped the PC Engines APU board with an Yepkit self-powered USB hub [25], three USB-based CAT4 MF910 MiFis [27] and one WiFi card [4]. The reason we chose the MF910 MiFi is because, at the time we selected the hardware, it was the most modern device sold by operators in the four countries where nodes operate.

In the prototype validation phase, this implementation presented some major obstacles. While the APUs proved to be very stable, the Mifis proved more challenging than expected. First of all, in the last quarter of 2016, the MiFis' vendor issued a forced update to the firmware. The update was applied despite the fact that we took special care to configure the devices not to receive automatic updates. As a result of the forced update, all our MiFis became inaccessible for the OpenMBB system. Furthermore, the MiFis themselves were prone to resets or to enter a working state (transparent PPP) from which we could only restore them to normal operation by draining their batteries, or perform a manual reboot by pushing the power button. Finally, after 6 months of operations, some of the MiFis showed clear signs of bloated batteries. This problem brought serious safety concerns for the nodes operating in places other than our own (controlled) premises (e.g., public transport vehicles). These obstacles forced us to modify the hardware configuration and use internal modems operating in the miniPCIe slots of the APU board.

Current Node Configuration: We decided to increase the control over the OpenMBB node and base its implementation on a dual-APU system. One of the two APUs in each node has two MC7455 miniPCI express (USB 3.0) modems [21], while the other has one MC7455 modem and a WiFi card. We chose Sierra Wireless MC7455 as our 4G modem since, at the time of the upgrade, it was supporting the most recent category (CAT6) an industrial grade modem could provide. This design eliminates the risk due to batteries, avoids any forced updates (the new modems are not routers), simplifies resets (no draining of battery) and increases our overall control over the system.

Takeaways: APUs showed very stable performance, while re-purposing the MiFis to behave as simple modems presented major challenges (e.g., forced updates and battery problems). We thus propose a more compact and robust node configuration that relies on internal modems operating in miniPCIe slots. This also simplifies the node since we avoid potential NAT and routing issues by avoiding MIFIs.

4. NODE SOFTWARE IMPLEMENTATION

In this section, we describe in detail the node software ecosystem and present the justification for our implementation choices.

4.1 Software Ecosystem

Figure 2 presents the different elements that coexist in the OpenMBB node software ecosystem. We divide the ecosystem into node management software, node maintenance software and experimentation enablers.

The **node management software** integrates a set of core components that run continuously in the background. They perform low-level work in line with the normal jobs expected on any mobile device or computer. These include (i) a *Device Listener*, which detects, configures and connects network devices, (ii) a *Routing Daemon*, which acquires an IP address through DHCP, sets up routing tables, (iii) a *Network Monitor*, which monitors interface state, checks the connectivity of the different interfaces and configures default routes. The node operates behind a firewall, which we configure with strict rules to increase node security.

The **node maintenance software** integrates components that monitor the node status and trigger actions to repair or reinstall when malfunctioning. A *system-wide watchdog* ensures that all core components (node management) are running. However, the watchdog alone was not enough to handle all different scenarios, especially mobility. During the first few months, we experienced loss of connection to nodes because of problems that watchdogs could not tackle, such as file system corruptions which can occur due to frequent

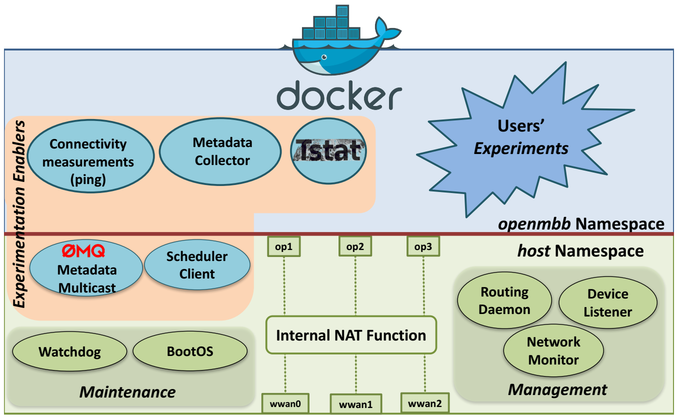

Figure 2: Node Software Ecosystem.

sudden power loss in mobile nodes. Thus, we defined and implemented a robust node recovery method, called BootOS, that enables a hard restart of the node (i.e., a reinstallation of the operating system to a known working baseline). This method allows us to recover both from file system errors that prevent system boot-ups, and also due to software configurations that may lead to loss of connectivity. To achieve this goal, we trigger a two-stage boot loader process at node start-up. In the first stage, we start the BootOS, which resides entirely in RAM and only uses read-only hard-drive access for its normal operation. The BootOS verifies that the filesystem of the APU is not corrupt, and that no forced reinstallation has been requested. It then proceeds to boot the MainOS, which contains the OpenMBB system software. If the filesystem is corrupt, or in case of a forced reinstallation, the BootOS reinstalls an image of a known working installation.

The **experimentation enablers** include the scheduling client, the default experiments, and the services for external experiments. Within the node software ecosystem, we differentiate between the user experiments, and the host management and maintenance software. We implemented this separation by configuring a *separate network namespace* (the `openmbb` namespace) where experiments run. This increases our control over the ecosystem and limits the impact external users can have on the node. This separation further allows us to account (as part of the scheduling system) the traffic volume each user consumes. We require that each experiment runs inside a virtualized environment (Docker container) to ensure separation and containment of processes. The *Scheduling Client* communicates with the Scheduler to enable experiment deployment per user request. It periodically checks for new experiment containers to run in the node and deploys them in advance to their scheduled execution time. Section 5 offers more details on the scheduling system. The *metadata broadcasting service* runs continuously in the background and relays metadata through ZeroMQ [26] in JSON [11] format to experiment containers. The nodes periodically run connectivity measurements (e.g. ping), and this together with metadata allow us to monitor the node's state and the overall health of the platform. Furthermore, *Tstat* [8] passive probe continuously provides insights on the traffic patterns at both the network and the transport levels, offering the experimenters additional information on the traffic each interface exchanged during an experiment.

Takeaways: Containment of users activity in the node is paramount to avoid security risks, node malfunctioning events, unreliable results and, more severely, node loss. We prevent foreign unauthorized access to the node with a strict firewall. Then, continuous monitoring of the platform is crucial and we enable it by implementing monitoring functions in the node management software. Node maintenance is expensive, so it is important to forge the node as a self-healing system. We implement this functionality in the node maintenance software that takes automatic actions when the node malfunctions.

4.2 Experiment Containment

Docker Virtualization. The node design we propose mandates that OpenMBB users execute their experiments inside Docker containers, which provide isolation from the host node. This is true both for default monitoring measurements and external users experiments. Docker containers are based on a layered file system, where a container can reuse layers shared with other containers.

OpenMBB provides the default base image for the experiment containers, which integrates the base operating system installation with default tools that are potentially useful for many experiments. The lightweight containers provide just the contents that are unique for the particular experiment, significantly reducing the download and deployment time overhead and accountable traffic volume. Experiments running inside a container have access to the experimental network interfaces. They can read and write on their own file system, overlaid over that of the base OpenMBB image. Finally, there are specific paths (e.g., `/OpenMBB/results/`) where the experiments can write their results and that the node automatically transfers to the OpenMBB servers. Our public software repositories contain all the files necessary to build new user experiments, as well as experiments templates and examples.

Internal NAT Function. To ensure the minimum impact of user experiments gone wrong, we define the *openmbb network namespace* where experiments containers run. For each physical interface that the network-listener detects as available, we create a virtualized ethernet, `veth`, interface pair, and move one end to the openmbb namespace. We then add routing rules in the network namespace to allow routing by interface. In order to allow the network devices in the host namespace to communicate with the ones in the openmbb network namespace, we define an internal Network Address Translation (NAT) function. We use `iptables` NAT masquerading rules in the host namespace to configure the NAT function. Finally, we add the corresponding routing rules to map each `veth` interface to the correct physical interface.

Overheads Quantification. The internal network design introduces two potential overheads that might impact the performance measurements, namely (i) the internal NAT function that connects the network devices in the host namespace with their corresponding duplicates in the `openmbb` namespace and (ii) the Docker containers we use to separate the processes that correspond to a certain experiment that runs inside the container. Thus, prior to detailing the measurement results of different commercial MBB operators, we focus here on these two design overheads and aim to quantify their impact (if any) on performance measurement results. More specifically, we quantify the delay overhead by running ICMP ping measurements, and the impact on throughput by running http downloads.

To instrument our system benchmarking measurements we use a single APU node running the Debian Stretch OpenMBB image with a local Fast Ethernet link. Using a local link allows us to minimize the impact of the network on our measurements, and focus on the impact of the system overheads. We run http download measurements with `curl` and ICMP ping measurements with `fping` to quantify the impact of the internal NAT function and of the Docker virtualization. We focus on four configurations for our testing setup, namely: no NAT and no Docker (experiments run in host namespace), no NAT but Docker (experiments run inside a Docker container in the host namespace), internal NAT and no Docker (experiments run in the `openmbb` namespace) and internal NAT and Docker (experiments run in side a Docker container in the openmbb namespace).

To quantify the delay overhead, we collect 1,000 RTT samples against the Google DNS server 8.8.8.8 on the ethernet connection on all four configurations. Figure 3 shows the results of the measurements. We conclude that the overhead of the NAT function internal to the node is insignificant. In average, we see a penalty in the order of 0.1ms, (i.e., in the range of clock granularity in Linux systems). We note that the Docker and NAT combination introduces a slight delay, which is not overwhelming.

For the throughput measurements, we download 1GB of data from a server we configure in the local network. We collect 30 samples for each testing configuration. In Figure 4 we show the cumulative distribution of download speed per namespace and operator, for each of the different targets. We find that there is a 1% perfor-

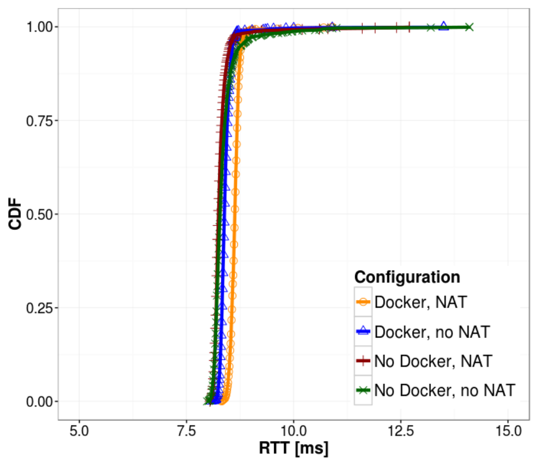

Figure 3: CDFs of ICMP RTTs [ms] measured against 8.8.8.8 per testing configuration over Fast Ethernet link.

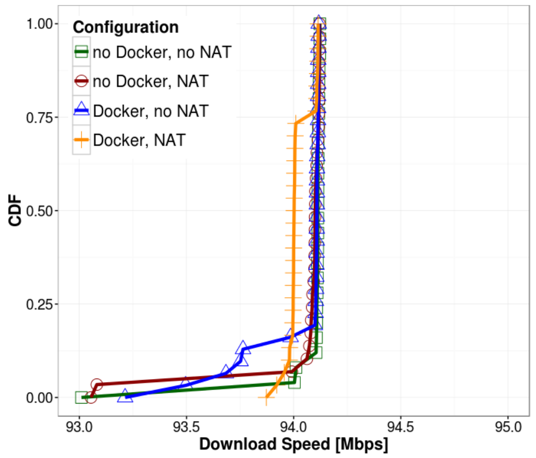

Figure 4: CDFs of Downloads Speed [Mbps] measured per testing configuration over Fast Ethernet link.

mance penalty that using the internal NAT function and the Docker virtualization introduces in average. We report no direct impact of using the Docker containers, which we expected, since the purpose of the Docker virtualization is purely for experiment containment.

Takeaways: Our priority in the node software implementation phase is keeping the nodes within normal functioning parameters for as long as possible and limiting direct maintenance intervention, while allowing external users to run a wide range of complex measurements with minimum interference. To achieve this, we separate the network namespace where users can run their experiments from the host namespace, where the monitoring and management software runs. This introduces two potential overheads in the system, which we quantify and show that have little or no impact.

5. USER ACCESS AND SCHEDULING

We provide access to the OpenMBB platform through a user-friendly interface consisting of an AngularJS-based web portal. As part of the OpenMBB federation with the Fed4FIRE [7] initiative, the user access follows the Fed4FIRE specifications in terms of authentication and resource provisioning. Through the portal, experimenters interact with the OpenMBB scheduler, which is in charge of deploying the experiments without allowing the users to directly access the nodes. The platform has a scheduler API that allows to bypass the web portal for experiment deployment automation. The scheduler ensures that there are no conflicts between users when running their experiments and assigns resources to each user based on their requirements and resource availability.

Given the challenging scenarios we aim to capture in our testbed, computing nodes in OpenMBB have potentially unreliable connectivity and low bandwidth. This is the norm for mobile buses, trains and trucks, which follow the schedule of the host vehicle. Experiment scheduling therefore accounts for two factors. First, the node may not have connectivity at the time of the experiment. Second, a high lead time when deploying containers means that experiments should be deployed early. However, experimenters may require to run synchronous measurements on multiple nodes. The common approach to task scheduling and decentralized computing, which deploys jobs to registered nodes based on their availability, struggles with these constraints. Therefore, for the OpenMBB scheduler, we follow a calendar-based approach, assigning time slots to experiments. Deployment of experiment containers takes place up to 24 hours in advance, as soon as the node retrieves information about the assigned task. This allows both immediate scheduling on nodes that are not otherwise occupied, and to schedule synchronous experiments on low availability nodes well in advance. It also allows to synchronize experiment runtime with vehicle schedules when available.

In addition to managing the time resource, the scheduler handles data quotas assigned by the contracts with the MBB operators. We assign to each experimenter a fix data quota, called user quota, which they can use for deploying and running their experiments. In addition, we may assign users a quota on computing time. We designed the quota system to provide fair usage of the available resources to experimenters. An important factor to ensure fairness in day-to-day usage, is that a certain data quota is reserved by the experimenter in advance, and subtracted from the user quota for the duration of the experiment. Experimenters may subsequently refund the remaining bandwidth. Hence, it is not possible to block large quantities of resources without having been assigned the necessary budget, even if the resources are not actually used.

From March 2016 until March 2017, the OpenMBB scheduler has been actively used by 30 users. A total of $75,002$ experiments have successfully ran on the platform, while $7,972$ scheduled experiments failed. There are many different reasons for failed experiments, for example that the container exits unexpectedly or the data quota is exceeded. Note that these failures are expected especially for the new users that are trying to familiarize themselves with the platform. We are running an open conversation with our users, gathering feedback from them and updating the user access and scheduling policies accordingly.

Takeaways: Resource allocation and experiment scheduling on OpenMBB is challenging because nodes have potentially unreliable connectivity (e.g., nodes in mobility scenarios) and limited data quota due to commercial-grade subscriptions. A calendar-based approach for scheduling addresses these requirements by taking into account per user and per node data quota, and synchronized experiment start time.

6. OPEN EXPERIMENTATION

Since the platform design phase, we have been working together with our users to understand their requirements from the system and which experiments have the highest appeal. We group these experiment in three categories that our users are currently curating and have been already actively using. The distribution of experiment runs to date among these categories is: Mobile Broadband Performance (19%), Service Oriented QoE (%36) and Innovative Protocols and Services (%45). The volume of data experiments in different categories varies, with Service Oriented QoE taking the largest quota (%60), while Innovative Protocols and Services are the least demanding (%10), despite registering the largest number of experiment runs. We further detail each category and provide examples of experiments and analysis we can perform using OpenMBB.

6.1 Mobile Broadband Performance

To measure a mobile network in a reliable and fair way, it is important to identify the metrics that accurately capture its performance. Different stakeholders have different metrics of interest and we argue that OpenMBB is able to cater to all of them. For example, regulators need connectivity, coverage and speed information to monitor whether operators meet their advertised services. Operators are interested in spatio-temporal data reporting the operational connectivity information to further identify instability and anomalies.

One important feature of the OpenMBB platform is that its deployment in public transportation vehicles allows users to evaluate MBB performance in diverse and complex urban mobility environments. A unique characteristic of this deployment is the repeatability of measurements obtained by many runs on the same route, at different hours. For example, Figure 5 shows RTT (ICMP ping) measurements for an operator in Sweden, as measured by the node operating aboard the same bus during several working days. In the figure, dot colors encode the range of values for the measured RTTs and we observe variations in RTT among different trips through the same location. Repeated measurements provide high confidence and diminish noise in the data, whereas measurement samples at the same location but at different hours allow for the analysis on the time-of-the-day effect (e.g., rush hour versus normal hours).

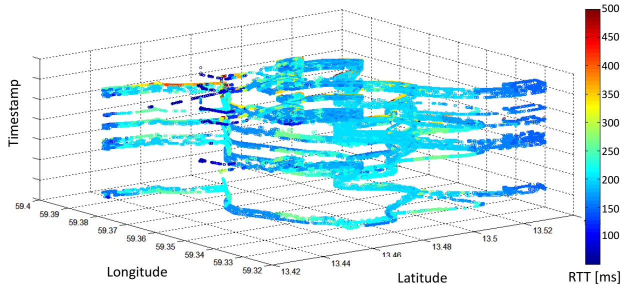

Figure 5: 3D graph average RTT for an operator in Sweden. Multiple laps are shown using the Y-axis offset based on relative timestamps to visually show the different trips.

6.2 Service Oriented Quality of Experience

An important measurement dimension to explore comes from the great interest in how users perceive individual services and applications over different terminals (e.g., mobile phones, tablets, and computers). The recent proliferation of user-centric measurement tools (such as Netalyzr [12]) to complement available network centric measurements validates the increasing interest in integrating the end user layer in network performance optimization. OpenMBB enables experimentation with essential services and applications, including video streaming, web browsing, real-time voice and video, and file transfer services. The service oriented measurements give a good bases for investigating the mapping from Quality of Service to Quality of Experience. With such a mapping, operators can gain better understanding of how their customers perceive the services delivered by their network. From the end users and service providers perspective, they could acquire more knowledge of the performance over different MBBs and then choose the network that delivers the best quality for services that are of interest to them. Furthermore, application developers (e.g. Youtube, Netflix and Spotify) heavily rely on the underlying network characteristics while optimizing their services for the best user's experience.

To showcase the capabilities of the platform, Figure 6 reports on the web page load time (PLT) we measure in the OpenMBB platform using a headless browser to fetch two popular websites (www.bbc.com and www.ebay.com) from the nodes operating in four countries with OpenMBB coverage. If we focus on the PLT as an objective indicator for the quality of experience and track it in comparison with the rich metadata information, we further enable the analysis of the mapping between QoS metrics to the end-user experience.

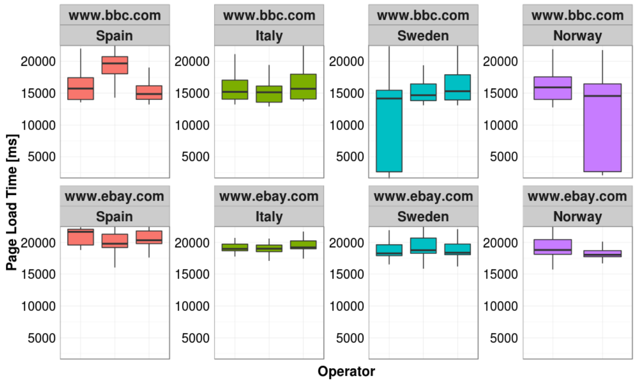

Figure 6: Page download time measured for 11 commercial operators using OpenMBB. We collect data from 37 nodes operating in 4 different countries (Spain, Italy, Sweden and Norway) while fetching two websites (`www.bbc.com` and `www.ebay.com`). Each sub-plot corresponds to a country-target pair, and each boxplot corresponds to a unique operator we measured in the corresponding country.

This is particularly interesting to dissect, since we observe large variation of the PLT for the same website in Sweden and Norway.

6.3 Innovative Protocols and Services

Another significant use case for OpenMBB is investigating the impact of middleboxes in the current Internet ecosystem. These range from address and port translators (NATs) to security devices to performance enhancing TCP proxies. Middleboxes are known to introduce a series of issues and hinder the evolution of protocols such as TCP. Since middleboxes are ubiquitous in MBB networks [22–24], OpenMBB offers an excellent opportunity to observe and characterize middlebox operations in the context of real world MBB deployments. OpenMBB further enables assessment of existing protocols, paving the way for protocol innovation.

As an example, we investigated whether the operators we measure with OpenMBB use Performance Enhancement Proxy (PEP) [2] to improve end-users' Quality-of-Experience. These proxies provide higher performance and faster error recovery [5, 10, 13]. We ran throughput tests (http downloads) from all the OpenMBB nodes on different ports against the same responder, where we also run an instance of Tstat on server. We then used the analysis logs of the Tstat on the nodes (client-side) and cross-compared them with the server-side logs to examine whether the proxy splits the TCP connection.

Table 1 shows the global view of operators. Yes, Yes*, and No in the table mean always, sometimes, or never. Third column indicates the usage of the NAT in the operator network. For instance, op2 in Italy is always using NAT, and sometimes connections are routed through a PEP. Contrariwise, op1 sometimes assigns public IP addresses, but the HTTP traffic always goes through a PEP device. The forth column indicates if the performance seen on the client and server side are

	OP	NAT	PEP	# IP	L4 mangling
IT	op0	No	Yes	262	80
	op1	Yes*	Yes	129	80,443,8080
	op2	Yes	Yes*	1484	No
ES	op0	Yes	No	272	No
	op1	Yes	No	244	No
	op2	No	Yes	-	80
SE	op0	Yes*	Yes*	1652	No
	op1	No	Yes	3486	No
	op2	No	Yes*	4679	No
NO	op0	No	Yes*	472	No
	op1	Yes*	Yes*	46	No

Table 1: The summary of the operators and their setting.

different. A mismatch hints for the presence of a PEP. The fifth column illustrates the number of public IPs we see in the server-side traces (i.e., the "size" of the PEP boxes). Last column shows if the PEP changes the TCP headers (e.g., removing/adding/changing options), and on which ports. Overall, the picture varies with different PEP configurations for different operators.

7. CONCLUSIONS

In this paper, we reported on our experience designing an open large-scale measurement platform for experimentation with commercial MBB networks. OpenMBB is a completely open system allowing authenticated users to deploy their own custom experiments and conduct their research in the wild. The platform is crucial to understand, validate and ultimately improve how current operational MBB networks perform towards providing guidelines to the design of future 5G architectures. We described our experience with the OpenMBB system implementation and detailed the hardware selection for the OpenMBB measurement node, its software ecosystem and the user access and scheduling solution. We emphasized the versatility of the design we propose, both for the overall platform and, more specifically, for the measurement nodes. In fact, the node software design is compatible with a number of different hardware implementations, given that it can run on any Linux-compatible multihomed system. Our current hardware solution is the most fitting for the set of requirements and the predicted usage of OpenMBB, which we evaluated based on our discussions and interaction with the platform's users.

References

[1] Beaglebone. http://beagleboard.org.

[2] Border, J., Kojo, M., Griner, J., Montenegro, G., and Shelby, Z. Performance enhancing proxies intended to mitigate link-related degradations.

[3] CAIDA-ARK. http://www.caida.org/projects/ark/.

[4] Compex. WLE600VX: http://www.pcengines.ch/wle600vx.htm.

[5] Farkas, V., Héder, B., and Nováczki, S. A Split Connection TCP Proxy in LTE Networks. In *Information and Communication Technologies* (2012), Springer, pp. 263–274.

[6] FCC. 2013 Measuring Broadband America February Report. Tech. rep., FCC's Office of Engineering and Technology and Consumer and Governmental Affairs Bureau, 2013.

[7] FED4FIRE. http://www.fed4fire.eu/.

[8] Finamore, A., Mellia, M., Meo, M., Munafo, M. M., Torino, P. D., and Rossi, D. Experiences of internet traffic monitoring with tstat. *IEEE Network 25*, 3 (May 2011), 8–14.

[9] Hirth, M., Hossfeld, T., Mellia, M., Schwartz, C., and Lehrieder, F. Crowdsourced network measurements: Benefits and best practices. *Computer Networks 90* (2015), 85 – 98.

[10] Ivanovich, M., Bickerdike, P. W., and Li, J. C. On tcp performance enhancing proxies in a wireless environment. *IEEE Communications Magazine 46*, 9 (September 2008), 76–83.

[11] JSON. http://www.json.org/.

[12] Kreibich, C., Weaver, N., Nechaev, B., and Paxson, V. Netalyzr: illuminating the edge network. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement* (2010), ACM, pp. 246–259.

[13] Necker, M. C., Scharf, M., and Weber, A. *Performance of Different Proxy Concepts in UMTS Networks*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005, pp. 36–51.

[14] networld2020. Service level awareness and open multiservice internetworking - principles and potentials of an evolved internet ecosystem, 2016.

[15] Odroid. http://www.hardkernel.com.

[16] OOKLA. http://www.speedtest.net/.

[17] PcEngines. APU: https://www.pcengines.ch/apu2c4.htm.

[18] Planetlab. https://www.planet-lab.org/.

[19] RaspberryPi. http://www.raspberrypi.org.

[20] RIPE. https://atlas.ripe.net/.

[21] Sierra-Wireless. MC7455 miniPCI express (USB 3.0) modem: https://www.sierrawireless.com/products-and-solutions/embedded-solutions/products/mc7455/.

[22] Vallina-Rodriguez, N., Sundaresan, S., Kreibich, C., Weaver, N., and Paxson, V. Beyond the radio: Illuminating the higher layers of mobile networks. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services* (2015), ACM, pp. 375–387.

[23] Wang, Z., Qian, Z., Xu, Q., Mao, Z., and Zhang, M. An untold story of middleboxes in cellular networks. In *Proc. of SIGCOMM* (2011).

[24] Xu, X., Jiang, Y., Flach, T., Katz-Bassett, E., Choffnes, D., and Govindan, R. Investigating Transparent Web Proxies in Cellular Networks. In *Proc. of Passive and Active Measurement* (2015).

[25] Yepkit. USB Switchable Hub YKUSH Homepage: https://www.yepkit.com/products/ykush.

[26] ZeroMQ. http://zeromq.org/.

[27] ZTE. USB-based CAT4 MF910 MiFi, product specification: http://www.ztemobiles.com.au/downloads/User_guides/MF910_Help1.0.pdf.

Disclaimer