MONROE

Measuring Mobile Broadband Networks in Europe

H2020-ICT-11-2014
Project number: 644399

Deliverable D1.1

Report on use cases

Editor(s): Min Xie, Andra Lutu and Özgü Alay
Contributor(s): Håkon Lønsethagen, Anna Brunstrom, Audun Fosselie Hansen, Vincenzo Mancuso and Marco Mellia

Work Package: 1 / System Design and Software Implementation
Revision: 0.1
Date: September 15, 2015
Deliverable type: R (Report)
Dissemination level: Public

European Commission
Horizon 2020
European Union funding
for Research & Innovation

Abstract

In this report, we describe representative use cases for the MONROE platform. The 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 examples of these use case categories. 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. The requirements of these experiments will serve as an input to the design of the MONROE system, and will ensure that the platform can cater for the requirements of the use cases, hence, fulfill the expectations of the stakeholders.

Participant organisation name	**Short name**
SIMULA RESEARCH LABORATORY AS *(Coordinator)*	SRL
CELERWAY COMMUNICATION AS	Celerway
TELENOR ASA	Telenor
NETTET SVERIGE AB	NET1
NEXTWORKS	NXW
FUNDACION IMDEA NETWORKS	IMDEA
KARLSTADS UNIVERSITET	KaU
POLITECNICO DI TORINO	POLITO

1 Introduction

Mobile Broadband (MBB) networks are becoming increasingly important in citizens' everyday life as well as for business applications. The use of MBB networks has exploded over the last few years due to the immense popularity of mobile devices such as smartphones and tablets, combined with the availability of high-capacity 3G/4G mobile networks. Emerging applications running on MBB networks are becoming more advanced and demanding. Given the increasing importance of MBB networks and the enormous expected growth in mobile traffic, there is a strong need for a better understanding of the fundamental characteristics of MBB networks and their relationship with the performance parameters of popular applications. To this end, we observe an increasing attention and interest towards MBB measurements. However, in order to meet different needs of a diverse set of stakeholders in the MBB ecosystem, it is essential to harmonize the metrics to characterize MBB networks, the measurement techniques to obtain these metrics and the approaches used for reporting these metrics.

MONROE's vision is to provide a flexible, open and industry-grade measurement platform for all its stakeholders to meet their diverse interests and needs. To this end, MONROE will design and run experiments in order to identify key performance and reliability metrics of MBB networks. The data collected by the platform will be made open. Moreover, a selected subset of the designed experiments will be provided as Experiments as a Service (EaaS) to interested parties, as agreed by the consortium. Furthermore, MONROE will be open to external users allowing them to deploy their custom experiments on the platform. With this, our ambition is to pave the way for assessing and improving the user experience for services that are running on the current mobile broadband infrastructure, while also providing feedback to upcoming 5G technologies.

MONROE platform aims to fulfill the expectations of different stakeholders. For MBB operators, it is crucial to track the performance of their networks from the end users' perspective. Results from MONROE will help operators with more accurate frequency planning, more cost-efficient investments, and better network utilization. Transport/logistic companies and emergency services depend heavily on the quality of mobile networks in their operations, and can benefit directly from measurements that investigates the reliability of these networks and methods that focus on improving it. Researchers and innovators needs an experimental platform in order to evaluate the performance of novel protocols and services in a real operational setting. In order to guide regulations and spur competition, regulators and society at large require large-scale independent measurements for assessing the stability and performance of MBB networks. Consumers need objective information on the cost, performance and reliability balance of different operators in order to make informed choices on which network provider to choose. In this document, we target to translate these expectations into representative use cases for MONROE and the experiments that enable these use cases. Throughout the testbed design phase, we review and analyze a wide range of these use cases in order to ensure the flexibility that is necessary for supporting different experiments that the end-users may introduce.

We identify three main categories of uses cases: (i) key mobile broadband metrics, (ii) application performance and (iii) protocol and service innovations. These three categories include several examples of use cases that are selected based on a thorough review of the state-of-the-art and a close interaction with academic and industry partners. We further identify a set of experiments in order to enable these use cases. The experiments described in this document will be carried out by the consortium and the corresponding open data and selected experiments (i.e. provided as EaaS) will be made public. In other words, different external users can deploy readily available experiments on the MONROE platform, in order to collect their own dataset, in addition to accessing to MONROE's open dataset.

The purpose of the first use-case category is to provide a rich dataset of key performance metrics which analyzes the MBB providers in terms of reliability and performance. The set of key performance metrics, together with the metadata will be molded into a network tomography tool that can enable further analysis which may be of interest to stakeholders. The second use case category focuses on the application performance and considers popular applications such as web and video services. The final use case category is dedicated to improving the MBB networks through novel protocols and services. The use cases here involve assessment of the paths, assessment of the protocols as well as leveraging multihoming capabilities of MONROE.

In order to properly provision the capabilities of the MONROE platform, we also consider several complementary and service specific use cases. These use cases are of interest to different stakeholders and they will be accommodated by the MONROE platform. However, the corresponding experiments will not be carried out by the MONROE consortium. The examples of these use cases are described in Appendix A. In Appendix B we list the acronyms we use in this document.

2 State of the Art

During the past years, we have seen increased interest in the networking community from different parties (e.g., researchers, operators, regulators, policy makers) in measuring the performance of mobile broadband networks. In this section, we aim to provide a condensed comprehensive review of some of the most relevant approaches that strive to shed light on the mobile broadband ecosystem. The use cases we identify in this document are in part the results of surveying the existing mobile broadband offerings.

In order to cater the need for open large-scale MBB measurements and to address the scarcity of available measurement platforms, several crowdsourcing approaches emerged over the past years, either from the research environment, e.g., Netalyzr [12], NetPiculet [26], or commercial-oriented, e.g., OpenSignal [18], RootMetrics [19] or MobiPerf [1]. These approaches leverage the wide adoption of mobile devices in the world and depend on the willingness of end-users to run the proposed tests. We note that the common vision of these tools is to identify and monitor a set of significant metrics which can accurately describe mobile broadband performance to the interested parties. For example, commercial-oriented OpenSignal proposes a holistic approach for building MBB coverage maps by retrieving the connectivity-related metadata from user devices. They introduce the notion of "time coverage" which provides statistics for the time a device has been using a certain radio access technology in order to provide the end-user the possibility to make informed decisions in terms of the preferred MBB provider in a certain area. Similarly, RootMetrics defines a set of key performance metrics which allows for network benchmarking, with the intent of rating different providers available in a certain geographical area. Additionally, tools such as NetPiculet or Netalyzr aim to shed light on the infrastructure of broadband providers with the purpose of informing protocol and application design.

There are several research projects [12–16] that use custom-designed apps to crowdsource and measure the performance of MBB providers and popular Internet applications, with a main focus on web browsing [27] and video streaming [25]. For example, MobiPerf [1] enables mobile network performance analysis [16]. The app builds on top of the Mobilyzer open library [17] and tracks a series of network performance metrics, such as HTTP benchmark downloading latency and bandwidth, traceroute with latency to different hops, ping latency, DNS lookup latency, TCP uplink and downlink throughput or RRC states metrics. Other similar relevant measurement efforts from the research community include [22–24].

With the increasing popularity of web and video-related services over MBB networks [9], there is a magnitude of research studies that focus on understanding the correlation between the network quality of service (QoS) metrics and the quality of experience (QoE) of the end-users [3, 7, 9]. In particular, this is appealing to operators, who continuously strive to provide the best service to their subscribers in order to increase their customer base. At the same time, the end-users themselves are looking for relevant metrics that can objectively assess the performance of popular applications over different MBB providers.

Even more, alongside the attention coming from end-users, businesses or operators, there is rising interest from regulators for defining and monitoring a representative and unitary set of metrics that accurately captures the performance of today's broadband services in practice. In this sense, several of them (e.g., FCC, Ofcom and Anatel) have translated these efforts into national projects in collaboration with commercial partners such as SamKnows [20], which specializes in home and mobile broadband performance evaluation. However, in order to allow for an open an unitary approach as well as the comparability of measurements, a common open framework is needed. This has been hard to achieve due to the proprietary nature of the measurement efforts, as is the case of [18–20], making it difficult for regulators to view measurement results from a harmonized and macroscopic scale. In this sense, several open measurement methodologies [2, 21] have been proposed with the goal of supporting the creation of inter-operable large-scale testbeds and advance a common approach on network performance characterization. The Internet Engineering Task Force (IETF) Large-Scale Measurement of Broadband Performance (LMAP) is currently working towards standardizing an overall framework for large-scale measurement platforms.

The MONROE platform complements the existing experimental platforms by providing unique features in the field of network-controlled mobile measurements. Three key aspects of MONROE that makes the platform unique are: repeatability and controllability of measurements for precise and scientifically verifiable results (even for the mobile scenarios), support for demanding applications such as web and video services and support for protocol and service innovation. These aspects sets up MONROE in an excellent position to advance the state-of-the-art measurement tools and platforms. Based on the literature survey we perform and considering the unique features of MONROE, we shape the main categories of use cases that we present in this document. We divide the use-cases into three main categories, focusing on: identifying and monitoring key MBB performance parameters, assessing the application performance and facilitating protocol and service innovation.

3 Use Cases

In this section, we expand on the set of representative MBB use-cases that are supported by the MONROE platform and the experiments that will be implemented by the consortium in order to enable these use-cases. Based on our analysis of the current MBB ecosystem, we identify three main categories for the use-cases, as follows:

3.1 Key Mobile Broadband Metrics

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, it is crucial to identify the metrics that accurately capture the performance and the conditions under which we evaluate these metrics. Different stakeholders have different requirements on the metrics supported by the MONROE platform. For example, on the one hand, regulators need connectivity, coverage and speed information collected from a third-party, independent platform to monitor whether operators meet their advertised services, and as a baseline for designing regulatory policies. On the other hand, operators are interested in time series reporting of operational connectivity data to identify instability and anomalies. Furthermore, application developers need to cross-check QoS parameters against the behavior of the underlying network to design robust services and protocols. From the above considerations, it is clear that the collection of data cannot be limited to transmission and packet-level statistics, but there is an obvious need for rich metadata to be associated with the performance and reliability measurements.

Next, we describe the set of experiments whose results are essential to characterize and monitor performance and reliability of MBB networks.

Key performance parameters: The metrics to be monitored include network-level QoS parameters such as UDP latency (Round trip time (RTT)) and packet loss, ICMP latency and packet loss, jitter (delay variation), UDP latency and packet loss under load, UDP loss runs (disconnections) and connection availability statistics. These provide a common and consistent view on the MBB network performance and further enable network benchmarking.

Network Metadata: The network metadata enables MONROE to capture the network context under which we measure the key performance metrics. The parameters we report in this category include but not limited to provider name, radio access technology (RAT) type, RAT-specific parameters (e.g., RSRP, RSRQ, RSSI) and network connectivity status. Network metadata is crucial not only for coverage information but also during the analysis of the measurements in order to understand the underlying factors that affect the performance.

Traffic analysis with Tstat: Traffic statistics, obtained through data correlation between incoming and outgoing traffic, can give reliable estimates of the network performance also from the user perspective. Therefore, traffic analysis tools that monitor network traffic and report live traffic statistics are necessary. Relying on the POLITO traffic analysis tool, tstat [6], the measured data is ready to be post-processed at the flow-level and visualized as either time plots or aggregated plots over different time scales. Tstat computes over 100 different performance statistics at both the IP and TCP layers, allowing a good insight into overall network performance.

3.2 Application Performance Measurement

With the development of performance measurement techniques, more users become interested in user-centric and application-oriented measurements that complement present network-centric measurement. The application-level performance measurements provided by MONROE provide a potential opportunity to explore user experience. Nowadays all players in MBB, especially the operators, care about quality of experience (QoE) and shift the focus to users. However, QoE is a complex concept integrating user perception and expectations, network and application QoS, and context data. It is both time- and resource-consuming to measure QoE subjectively. A more realistic and preferred approach for operators is to measure objective QoS and map the measurement to perceived QoE. MONROE could generate a sufficiently large amount of measurement data in a wide area to establish, test and verify a QoS-QoE mapping model for different services. This way, 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.

Considering that user perception varies from service to service, MONROE builds on top of the key MBB measurements while measuring application performance (i.e., for each application, additional performance metrics are defined). Based on the mobile data traffic trends and analysis of the user demand and system requirements, the following services are selected to be implemented by the MONROE consortium:

Video services: According to Cisco and Ericsson [4, 5], mobile video traffic is the major growth factor, accounting for more than 70% of the overall mobile traffic by 2019. Delivering video traffic not only consumes more resource but also needs to meet higher quality requirements. It is critical to understand how well MBB transmits video services to end users. Two examples of video services are considered in MONROE: video streaming and video surveillance.

Video streaming is a dominant data service in both current and emerging mobile markets and is anticipated to be the dominant service also in the coming years. However, even today's 4G networks can not always guarantee smooth delivery of video streaming. Therefore, it is important for operators to know when network performance would be intolerable for users and make users churn. From the end user perspective, video is the dominant service for most of the users and they demand high quality video streaming especially when they are mobile. The content providers of video streaming services, such as Youtube and Netflix, use HTTP-DASH widely. Therefore, MONROE focuses on the performance of HTTP-DASH over MBB networks, covering static and mobile scenarios.

Video surveillance systems integrated in mobile vehicles such as busses or trams produce a large amount of video that cannot be delivered to the Security Center in real time. A practical solution is to store videos in the mobile vehicle and then upload videos to the central security system when the vehicle enters a depot with cheap and high speed wireless connectivity. However, this forces the Security Operator to wait for several hours to access the video in case they need it. In the context of the MONROE platform, we can experiment a novel solution which would provide the Security Operator with on-demand access to the video feed, by using a console which allows her to select the time interval of the video she would like to see, so that a on-demand video upload can be triggered and executed. The MONROE system will be used to evaluate the efficiency of the video surveillance system to upload video over MBB, with consideration of the video size, the upload rate, the cost of each interface and the deadlines specified by the Security Operator.

Web traffic: Web services are prevalent over Internet and are a big contributor to current MBB traffic, particularly in the emerging HTTP/2 standard. Therefore, it is utmost important to assess the performance of web services over MBB networks considering both HTTP/1.1 and HTTP/2, especially for web service developers. The application-level performance metrics to be measured may be transaction time for web browsing (from the time the first connection is initiated and until the complete website is done) and latency under load for file transfer and web browsing [10].

Background traffic: On mobile devices, background traffic (*e.g.*, system update or peer to peer traffic) frequently occurs and may affect the perceived performance of other applications. It is interesting for service providers and mobile app developers to obtain the measurements on how the foreground traffic is correlated with the background traffic and how their performance is interrelated. This is especially interesting for delay sensitive applications since MBB networks are known to have large buffers, a phenomenon called bufferbloat [11], that can impair the user experience for these applications especially when there is background traffic.

3.3 Innovative protocols and services

MONROE not only provides measurement for existing protocols and services, but also allows for the flexibility of testing and assessing innovative protocols and services. MONROE deploys nodes in multiple countries, provides multihoming capabilities and supports stationary and mobile scenarios. The multi-homing aspect of MONROE enables researchers and business (e.g. Wifi service provide on trains and mobile video broadcaster) to experiment with protocols and algorithms that exploit multiple connections opportunistically, e.g., in parallel or by picking the one with the best available service to increase robustness and performance, or to achieve the best cost-performance ratio. Operators can benefit from the traffic offloading support between MBB and Wifi service in order to ease congestion and reduce the traffic load. A measurement testbed is necessary for them to assess the applicability, performance, and efficiency of different traffic offloading mechanisms in a practical environment. End users offload data mainly for cost control and/or use of better networks and algorithms that improves the end users experience is very much needed. End users are also subject to many issues introduced by different types of middleboxes (e.g. address and port translators (NATs), security devices and performance-enhancing TCP proxies) that impact the user's experience. Therefore, it is essential to have a platform to observe and characterize middlebox operation in real-world deployments in MBB networks.

Path support with active measurements: 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. It thus demonstrably becomes problematic to extend core Internet protocols, limiting the opportunities for optimization. Consequently, it is of paramount importance to understand the interaction between the Internet path and protocol optimization solutions. For example, the IP and TCP functionality (*e.g.*, MPTCP, ECN, TCP Fast Open) of different paths can be tested through active measurements running on the MONROE platform. The system allows experimenters to measure wide diversity of paths and can help determine whether a proposed solution has the required support or functionality from the MBB ecosystem.

Protocol performance: Various transport protocols and transport protocol parameters can be tested in MONROE to evaluate their performance over MBB and WiFi. Especially TCP performance is of great interest since it is the predominant transport layer protocol used by applications on mobile networks. However, TCP was not designed for MBB networks and a recent study showed that many TCP connections significantly under-utilize the available bandwidth in LTE networks [9]. TCP has many different parameters including but not limited to congestion control mechanisms, receiver buffer settings, slow start and loss recovery that can individually or collectively impact the performance of different applications. Another important aspect to be considered is how TCP handles the handover among different MBB technologies as well as among different access technologies.

Multipath transport: Multi-path transport has shown to provide benefits, from bandwidth aggregation to increased robustness by taking advantage of the diversity of the underlying paths. In order to provide a good quality of service over multiple paths, mechanisms such as path management, packet scheduling and congestion control are essential. Considering the dynamic characteristics of MBB networks especially under mobility, and Wifi networks that has very different characteristics than MBB networks, multipath transport over MBB and Wifi networks becomes an especially challenging problem. A smartphone is a common example for such a scenario, since it usually provides two different network interfaces: Wifi and MBB. In such challenging scenarios, we need novel designs for robust and adaptive mechanisms that needs to be tested and verified in real networks. MONROE facilitates such validation processes.

Traffic offloading between MBB and WiFi: Due to the surge of mobile data traffic, traffic offloading is becoming a new industry segment, used by both end users and operators. End users offload data mainly for cost control and/or use of better networks. The operators offload traffic from cellular networks to WiFi in order to ease congestion and reduce the traffic load. A measurement testbed is necessary for them to assess the applicability, performance, and efficiency of different traffic offloading in a practical environment. The MONROE platform, particularly the mobile MONROE nodes, will produce data that could benefit operators with more efficient network planning and benefit end users with enhanced user experience.

4 Hardware and Software Requirements

The MONROE platform will be assessed and validated such that platform's software and hardware do not become a bottleneck that can disrupt or distract the performance of the network interfaces and the running measurements. The design of hardware and software is thus driven by the requirements of the use cases we identify in this document. Though the vision of MONROE is to provide a flexible platform with measurement nodes that can accommodate a myriad experiments, we identify an inter-dependability between the specifications of the use case requirements and the capabilities of the platform. In other words, the use-cases may need to be molded in order to suit some features, capability, and limitations of available hardware and software.

In terms of hardware requirements, the use cases demand a measurement node that is capable of handling even the most demanding applications, such as video streaming and web services. Therefore, we need a powerful node that can provide the necessary processing power for such demanding applications. Furthermore, the node should be equipped with multiple interfaces (e.g. MBB, WiFi and LAN interfaces) in order to evaluate different access technologies. When benchmarking MBB networks, we need to keep in mind that different areas have different coverage profiles in terms of technology. Therefore, the selected modems should support the current LTE technology along with the previous 3G and 2G technology. Some of the envisioned experiments require testing in mobility scenarios, which further means that the nodes should include a GPS module in order to track the location of the nodes. We design the MONROE node to meet the above-mentioned requirements. Furthermore, to accommodate the needs of eventual external users, additional hardware components can be supported. For example, for measuring the energy efficiency of the protocols and services we can add energy probes to monitor the power supply unit.

In terms of software requirements, multiple experiments running on the node simultaneously need proper resource allocation and arbitration. Furthermore, the innovative services and protocols might require privileged access and kernel modifications (*e.g.*, test of path support and test of transport protocols). This requires a flexible and reliable software design, that can handle multiple experiments at the same time where each experiment can run in a different environment. Furthermore, different experiments might need to access metadata or GPS information and use it as input while running. Therefore, all the metadata should be provided to all experiments in a timely and secure manner.

In terms of the back-end system, MONROE consortium decided to visualize selected key MBB metrics in a user friendly website. The visualization tool will be near real-time. This imposes some additional challenges to the back-end system design, such as tight synchronization of the measurement data and metadata from the same node to the back-end database as well as the visualization system. Furthermore, node mobility will be supported and the varying locations of the mobile measurement nodes will be visualized on the map.

5 Conclusions

MONROE aims to provide a flexible, open and industry-grade measurement platform in order to meet the diverse expectations of all its stakeholders. The unique features of MONROE enables realistic and accurate monitoring of MBB performance and reliability as well as methods to improve them. In this document, through a comprehensive literature survey and numerous interactions with the stakeholders, we identify three main category of use cases to meet the demands of potential MONROE users: i) understand the basic performance of MBB; ii) comprehend how different applications and services are delivered by the MBB network to end users; iii) test and verify innovative protocols over MBB. For each use case category, we define representative use cases that are of interest to different stakeholders. We then establish a set of experiments to enable these use cases. These experiments will be implemented by the consortium and the corresponding open data along with EaaS for a subset of these experiments will be made open. MONROE's external users can define their own experiments on top of the MONROE platform and decide how they will utilize the MONROE platform. Several examples of such potential complimentary and service specific use cases are given in Appendix to demonstrate how MONROE could be leveraged at large.

References

[1] MobiPerf. http://www.mobiperf.com, 2014.

[2] Marcelo Bagnulo, Philip Eardley, Trevor Burbridge, Brian Trammell, and Rolf Winter. Standardizing Large-scale Measurement Platforms. *SIGCOMM Comput. Commun. Rev.*, 43, 2013.

[3] Pedro Casas, Michael Seufert, and Raimund Schatz. YOUQMON: A System for On-line Monitoring of YouTube QoE in Operational 3G Networks. *SIGMETRICS Perform. Eval. Rev.*, 41(2):44–46, August 2013.

[4] Cisco Systems, Inc. *Cisco visual networking index: Global mobile data traffic forecast update, 2014 - 2019*, February 2015.

[5] Ericsson. Ericsson mobility report: on the pulse of the networked society. http://www.ericsson.com/res/docs/2015/ericsson-mobility-report-june-2015.pdf. Accessed: 2015-09-10.

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

[7] F. Fund, Cong Wang, Yong Liu, T. Korakis, M. Zink, and S.S. Panwar. Performance of dash and webrtc video services for mobile users. In *Packet Video Workshop (PV), 2013 20th International*, pages 1–8, Dec 2013.

[8] IETF IP Performance Working Group. Ietf draft: Model based bulk performance metrics, Mar 2015.

[9] Junxian Huang, Feng Qian, Yihua Guo, Yuanyuan Zhou, Qiang Xu, Z. Morley Mao, Subhabrata Sen, and Oliver Spatscheck. An in-depth study of lte: Effect of network protocol and application behavior

on performance. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, pages 363–374, New York, NY, USA, 2013. ACM.

[10] Berkely ICSI. Icsi netalyzr: Network buffering test. http://netalyzr.icsi.berkeley.edu/info_buffer.html. Accessed: 2015-08-30.

[11] Haiqing Jiang, Yaogong Wang, Kyunghan Lee, and Injong Rhee. Tackling Bufferbloat in 3G/4G Networks. In *Proc. of IMC*, 2012.

[12] Christian Kreibich, Nicholas Weaver, Boris Nechaev, and Vern Paxson. Netalyzr: illuminating the edge network. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 246–259. ACM, 2010.

[13] Anh Le, Janus Varmarken, Simon Langhoff, Anastasia Shuba, Minas Gjoka, and Athina Markopoulou. Antmonitor: A system for monitoring from mobile devices. In *Proceedings of the 2015 ACM SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big (Internet) Data*, pages 15–20. ACM, 2015.

[14] Mahesh K Marina, Valentin Radu, and Konstantinos Balampekos. Impact of indoor-outdoor context on crowdsourcing based mobile coverage analysis. In *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges*, pages 45–50. ACM, 2015.

[15] Massimiliano Molinari, Mah-Rukh Fida, Mahesh K Marina, and Antonio Pescape. Spatial interpolation based cellular coverage prediction with crowdsourced measurements. In *Proceedings of the 2015 ACM SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big (Internet) Data*, pages 33–38. ACM, 2015.

[16] Ashkan Nikravesh, David R. Choffnes, Ethan Katz-Bassett, Z. Morley Mao, and Matt Welsh. Mobile Network Performance from User Devices: A Longitudinal, Multidimensional Analysis. In *Procs. of PAM*, 2014.

[17] Ashkan Nikravesh, Hongyi Yao, Shichang Xu, David Choffnes, and Z Morley Mao. Mobilyzer: An open platform for controllable mobile network measurements. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*, pages 389–404. ACM, 2015.

[18] OpenSignal™. *Open Signal Inc.*

[19] RootMetrics™. *RootMetrics.*

[20] SamKnows™. *Methodology and technical information relating to the SamKnows™ testing platform - SQ301-002-EN*, 2012.

[21] S. Sen, J. Yoon, J. Hare, J. Ormont, and S.Banerjee. Can they hear me now?: A case for a client-assisted approach to monitoring wide-area wireless networks. In *Proc. of IMC*, 2011.

[22] M Zubair Shafiq, Lusheng Ji, Alex X Liu, Jeffrey Pang, and Jia Wang. Characterizing geospatial dynamics of application usage in a 3g cellular data network. In *INFOCOM, 2012 Proceedings IEEE*, pages 1341–1349. IEEE, 2012.

[23] Joel Sommers and Paul Barford. Cell vs. WiFi: On the Performance of Metro Area Mobile Connections. In *Proc. of IMC*, 2012.

[24] Wee Lum Tan, Fung Lam, and Wing Cheong Lau. An empirical study on 3G network capacity and performance. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 1514–1522. IEEE, 2007.

[25] Florian Wamser, Michael Seufert, Pedro Casas, Ralf Irmer, Phuoc Tran-Gia, and Raimund Schatz. Yomoapp: A tool for analyzing qoe of youtube http adaptive streaming in mobile networks. In *Networks and Communications (EuCNC), 2015 European Conference on*, pages 239–243. IEEE, 2015.

[26] Zhaoguang Wang, Zhiyun Qian, Qiang Xu, Zhuoqing Mao, and Ming Zhang. An untold story of middleboxes in cellular networks. In *Proc. of SIGCOMM*, 2011.

[27] Zhen Wang, Felix Xiaozhu Lin, Lin Zhong, and Mansoor Chishtie. How far can client-only solutions go for mobile browser speed? In *Proceedings of the 21st international conference on World Wide Web*, pages 31–40. ACM, 2012.

A Potential Complimentary and Service Specific Use Case

We list next other potential use cases for the MONROE platform, which exemplify interesting questions which the consortium or other external users can answer using MONROE.

- **Differentiated Services.** Anticipating the future MBB and Internet where best-effort traffic will flow side-by-side with specialized services, the MONROE measurements and reporting approaches will become even more important. The challenge is to extract comparative measurements results and reports regarding network QoS for these two main categories of traffic flows. The regulator will monitor the general quality level of best-effort traffic to ensure a reasonable quality level is ensured, and the MONROE techniques and reporting approach must balance the needs and the views of the different stakeholders.

- **Device-to-Device Communications.** D2D is a representative application of IoT. Measuring D2D-related service performance is interesting for future service study. Its measurement may cover D2D pair discovery and connection, and exchange of social and context-aware data, tested on mobile and (optionally) static nodes. The D2D use case can be also applied to test new technologies like WiFi Direct connectivity, context-aware social networking, data connectivity in social networks, WiFi-aware discovery functionality, and Floating Content services.

- **Route analytics with network topology inference and analysis.** The generated network-level interconnection graph from different data sets will show how networks are interconnected and how traffic flows in the Internet. It is relevant to understand circuitous paths in the Internet that cause longer delays than expected. It also provides vital information for the net-neutrality debate. The measured metrics include path dynamics such as traceroute/mtr, Internet routing data, and prefix geolocation.

- **Channel quality modeling.** Channel models could characterize time correlation, spatial correlation (for nearby devices on the same node or neighboring nodes) and WiFi transmission opportunities by passively observing channel dynamics. This experiment-driven modeling may lead to more spectral-efficient design of user applications by exploring spatial and temporal correlations.

- **From big data to big questions.** The MONROE platform will collect a huge amount of measurement data, which can be used to discover interesting *research questions*. The first type of questions is related to identification of unknown problems on the network, protocols, or applications. Mathematical approaches can be proposed to classify the network elements (hardware and software) based on their importance and nescience to the identified problem. Then a quantification study can be conducted to investigate whether it is worthy to resolve the problem. Similarly, a mathematical analysis can be done towards the used tools. Then combining the tools and problems could lead to another new research questions.

  The second type of questions is to explore and expand available measurements. Example questions are: what should be measured? How well the model is understood? How to apply new methods of machine learning to find new models? Note that the intention is to find the right "‘questions"’ rather than the new/right models.

- **Video conferencing services**. Voice and video calls over Internet (VVoIP) are becoming an important tool to connect people, both professionally and socially. The real-time requirement and interactive communication mode make it more challenging to study the QoE of VVoIP than video streaming. WEB-RTC is one of the enabling technologies to test VVoIP over MBB networks. Considering that WebRTC

API provides a set of statistics from the application perspective (by browsers), the synthesis tool can combine WebRTC API statistics with network QoS measurement to expand the QoE understanding further into the end-to-end aspect.

- **Online gaming**, which has continuously grown recently, especially with fast growing mobile penetrations. Compared to other services, online gaming is more delay-sensitve and has stricter QoS/QoE requirements. To evaluate the impact of delay on user experience, Real Time Strategy and Card games (available on Facebook) are considered as examples.
- **Energy profile of MBB technologies and energy aware transport.** In the context of the ongoing deployment of 4G/LTE networks, there is increasing interest to better understand its performance and power characteristics, compared with 3G/WiFi networks under realistic models. Besides higher bit rate, lower latency and many other improvements offered by LTE, user equipment (UE) power saving is an important issue, which can be characterized and further optimized. The configuration of LTE-specific parameters implies a tradeoff between UE power saving, channel scheduling delay, and signaling overhead. Understanding this tradeoff is of further importance from the point of view of improving the current 3G/4G technologies as well as informing the design of the upcoming 5G technologies.

B List of Acronyms

MONROE Measuring Mobile Broadband in Europe

MBB Mobile Broadband

EaaS Experiments as a Service

IETF Internet Engineering Task Force

LMAP Large-Scale Measurement of Broadband Performance

HTTP Hypertext Transfer Protocol

DNS Domain Name System

QoS Quality of Service

QoE Quality of Experience

UDP User Datagram Protocol

TCP Transmission Control Protocol

RTT Round Trip Time

ICMP Internet Control Message Protocol

IP Internet Protocol

LAN Local Area Network

RAT Radio Access Technology

RRC Radio Resource Control

RSRP Reference Signal Receive Power

RSRQ Reference Signal Receive Quality

RSSI Received Signal Strength Indicator

GPS Global Positioning System

DASH Dynamic Adaptive Streaming over HTTP

NAT Network Address Translation

MPTCP Multi-Path Transmission Control Protocol

ECN Explicit Congestion Notification

5G Fifth Generation Mobile Networks

LTE (4G-LTE) Long Term Evolution

3G Third Generation of Mobile Telecommunications Technology

2G Second Generation Wireless Telephone Technology

Disclaimer