<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *The 4th International Forum on Research and Technologies for Society and Industry (IEEE RTSI 2018)*.

Citation for the original published paper:

Özgü, A., Mancuso, V., Brunström, A., Alfredsson, S., Mellia, M. et al. (2018)
End to End 5G Measurements with MONROE: Challenges and Opportunities
In:

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kau:diva-69299>

End to End 5G Measurements with MONROE: Challenges and Opportunities

Özgü Alay¹, Vincenzo Mancuso², Anna Brunstrom³, Stefan Alfredsson³, Marco Mellia⁴, Giacomo Bernini⁵, Hakon Lonsethagen⁶

¹Simula Metropolitan Center for Digital Engineering, Norway, ²IMDEA Networks Institute, Spain, ³Karlstad University, Sweden, ⁴Politecnico di Torino, Italy, ⁵Nextworks, Italy, ⁶Telenor Research, Norway

Abstract—To be able to support diverse requirements of massive number of connected devices while also ensuring good user experience, 5G networks will leverage multi-access technologies, deploy supporting operational mechanisms such as SDN and NFV, and require enhanced protocols and algorithms. For 4G networks, MONROE has been key to provide a common measurement platform and a set of methodologies available to the wider community. Such common grounds will become even more important and more challenging with 5G. In this paper, we elaborate on some key requirements for the design and implementation of 5G technologies and highlight the key challenges and needs for new solutions as seen in the context of 5G end-to-end measurements. We then discuss the opportunities that MONROE provides and more specifically, how a 5G-capable MONROE platform could facilitate these efforts.

I. INTRODUCTION

Mobile Broadband (MBB) networks underpin numerous vital operations of the modern society and are arguably becoming the most important piece of the modern communications infrastructure. 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 cellular networks. According to Cisco's Global Mobile DataTraffic Forecast [1], mobile data traffic grew 57% in 2017 and is expected to grow with an annual rate of 47% by 2021. Consequently, traffic from wireless and mobile devices will account for more than 63% of total Internet traffic by 2021. Given the increasing importance of MBB networks and the enormous expected growth in mobile traffic, there is a strong need to better understand the fundamental characteristics of MBB networks both in terms of (i) enhanced MBB use cases and capabilities as well as (ii) massive machine type communications and (iii) ultra-reliable and low-latency communications, which are the three dimensions identified by the International Telecommunication Union (ITU) for 2020 and beyond, as shown in Fig. 1 [2].

While the main advancement offered by 4G consisted in the improved general experience of the Internet access service, the 5G cellular network vision directs the attention into a multitude of specific use cases along the three above mentioned dimensions proposed by ITU. This also includes a rich set of actor roles, business relationships and sector specific ecosystems. Indeed, 5G is upcoming and is expected to bring in a number of novel features meant to enrich MBB services by

Fig. 1: IMT Vision Usage scenarios for 5G.

specifically targeting the so-called vertical sectors, as defined by the 5G Public-Private Partnership (5G PPP) [3]: e-Health, Factories of the future, Energy and Automotive. Enabling such verticals requires the introduction of high speeds, ultra-low latency, ultra-high reliability, massiveness of network scenarios, all while meeting the requirements of network and service sustainability. The solutions proposed by 5G PPP span from (i) New Radio (NR) techniques, which include the use of mmWaves, (ii) massive use of Software Defined Networking (SDN) and Network Function Virtualization (NFV) to make the network more flexible and adaptive, (iii) the introduction of computational elements close to the users (in the edge of the operator's network) regulated by network slicing techniques to support specialized services, (iv) the adoption of multiple and heterogeneous connectivity techniques, including (v) unconventional connectivity schemes, e.g., the use of Device-to-Device (D2D) communications, and (vi) the support for Internet of Things (IoT) architectures.

All the novel aspects introduced by 5G are complex and intertwined, and therefore hard to evaluate on paper. They require methodical experimental investigation, which is quite challenging because of the heterogeneity of the 5G scenario. In such context, the benefits of a common measurement platform and set of methodologies available to the wider community is key, as it has already been demonstrated by the Measuring Mobile Broadband Networks in Europe (MONROE) platform for the case of 4G [4]. The need for such common grounds will become even more important and more challenging with

Fig. 2: Main building blocks of the MONROE platform.

5G. MONROE offers the possibility to remotely and automatically orchestrate *containerized* experiments in fixed and mobile Linux machines connected to real operators' access networks, thus acting as regular customers, and generate rich metadata, which is key to interpret experimental results. As we show in this paper, extending MONROE to 5G offers a great opportunity to readily measure 5G features since early 5G releases, which would offer early feedback for better 5G service/protocol design. It would also require only limited effort integrating the new 5G radio capabilities in the measurement platform, using software containers and off-the-shelf USB/mPCIe hardware extensions. Specifically, in this paper, we follow-up on some key requirements for the design and implementation of 5G technologies, what are key challenges and needs for new solutions as seen in the context of measurements, and how an evolved platform for MONROE can facilitate these efforts.

The paper is structured as follows. Section II describes the MONROE platform and its current capabilities. Section III elaborates on the upcoming 5G paradigm and identifies key areas. Section IV presents the possibilities of conducting 5G-related research using the MONROE platform, as well as summarizes the main requirements for evolving the MONROE platform. Section V concludes the paper.

II. THE MONROE PLATFORM

MONROE project provides the first European transnational open platform for independent, multi-homed, large-scale measurements and experimentation in commercial MBB networks. The platform comprises both fixed and mobile vantage points (the MONROE nodes) in Norway, Sweden, Spain, Italy, Germany, Portugal, Greece and the UK. Each node connects to three mobile network operators (MNOs) as well as WiFi, and allows users to conduct a wide range of experiments, also under challenging mobility scenarios.

MONROE integrates two main components (Fig. 2): the *MONROE nodes* and, the *User Access and Scheduling System*.

A. Node Hardware and Software

Each MONROE node integrates 2 small programmable computers (PC Engines APU2 board) interfacing with 3

3G/4G MC7455 miniPCI express modems using LTE CAT6 and one WiFi transceiver. Nodes are multi-homed to 3 different Mobile Network Operators (MNOs). All software components used in the platform are open source and available online¹.

The software on the nodes is based on Debian GNU/Linux “stretch” distribution. All experiments run inside a virtualized environment (Docker container² with Linux) to ensure separation and containment of processes. MONROE enables both continuous monitoring measurements as well as external user experiments. Monitoring measurements include active measurements such as connectivity measurements (e.g., ping) and speedtest measurements [5] as well as the *Tstat* [6] passive probe that provides insights on the traffic patterns at both the network and the transport levels. Furthermore, to provide rich metadata to the experiment containers, the *metadata broadcasting service* runs continuously in the background and relays metadata through ZeroMQ³ in JavaScript Object Notation (JSON) format to experiment containers.

The “metadata subscriber” in each node listens to *all* the metadata topics, and the results are then transferred to the MONROE back-end servers. Once at the server, every data item is processed and stored in the MONROE database for public access. All experiment results are also synchronized with the MONROE servers and, as a final step, provided back to the experimenter through the web interface.

B. User Access and Scheduling

We provide access to the MONROE platform through a user-friendly interface consisting of an AngularJS-based web portal⁴. Through the portal, experimenters interact with the scheduler and deploy their experiments without accessing the nodes directly. The scheduler API is accessible to enable experiment deployment automation. The scheduler prevents conflicts between experiments (i.e., only one user can run an experiment on a certain node at a given time) and assigns resources to each user based on their requirements and resource availability. The scheduler offers the possibility to load pre-configured experiments, using customizable parameters. MONROE also offers a number of experiment templates, in the form of Docker containers that can be downloaded and modified. The result is a platform that offers Experiments as a Service (EaaS) to a wide spectrum of user profiles.

MONROE offers a unique opportunity for experimental validation of 5G applications and services. Next, we will discuss the key challenges and needs for new 5G solutions as seen in the context of measurements, and how an evolved MONROE platform could facilitate these efforts.

III. TOWARDS 5G: CHALLENGES

5G mobile networks are envisioned to support significantly faster mobile broadband speeds and increasingly extensive

¹<https://github.com/MONROE-PROJECT/>

²<https://www.docker.com/>

³ZeroMQ (ZMQ) distributed messaging: <http://zeromq.org>

⁴Web portals to the MONROE project and the experiment scheduler (registration needed for access): <https://www.monroe-project.eu/> and <https://www.monroe-system.eu/>

mobile data usage, as well as to enable the full potential of IoT. The applications and services of the future prioritize capacity, latency, reliability and efficiency differently as illustrated in Fig. 1. To support such diverse services, 5G is relying on technologies such as mmWave, small cells, the use of multi-access connectivity and network slicing enabled through the SDN and NFV paradigms, which leads to the challenges discussed in what follows.

A. Software-Defined Networking and Network Function Virtualization

5G will provide a scalable network infrastructure to meet the exponentially-increasing demands on mobile broadband access, both in terms of number of connected users and required bandwidth. Deep and extensive cloudification of services, with integration of edge and centralized clouds is a 5G key aspect that is accelerating the adoption of NFV and SDN technologies as key enablers of truly flexible and automated service management in 5G networks.

The evolution towards a 5G enabled service architecture is requiring the creation of a new service environment, built on top of a mesh of micro data centers, as much as possible distributed to cover the edge of network for end-user proximity, coordinated by advanced service management platforms. These platforms, including advanced virtual infrastructure management platforms enabling service components and functions to run as virtualized and containerized applications, are required to provide enhanced agility for creation and operation of new services based on NFV and SDN technologies. The migration towards fully softwarized networks and services running on distributed virtualized infrastructures managed, controlled and orchestrated with NFV and SDN tools will allow service providers to reduce significantly the operational impact of launching new services.

More than that, another key concept that will be introduced with 5G is network slicing, that is intended to enable operators to create multiple isolated logical network instances (i.e. slices) over the same underlying network infrastructure, each optimized according to specific service and business requirements of vertical industries (being considered as main consumers of network slices) and possibly owned by different tenants. 5G envisions network slices as composed by network functions, providing both user and control plane functions spanning across radio access, mobile edge and core network segments. Here, NFV and SDN technologies become key enablers for transforming 5G network slicing into reality. In particular, NFV and SDN are expected to provide the needed functionalities to automate the customized network slice resource allocation to fulfill the required quality of service and performance levels. Service management and orchestration platforms are required to evolve and leverage on NFV and SDN to deal with a wide and differentiated pool of resource types and thus coordinate the delivery of 5G end-to-end softwarized services composed by any combination of virtual, legacy, or SDN-based functions implementing the ground for isolated network slices. In this context, understanding and

assessing the end-to-end performance of the composed service offered by network slices constitute a key challenge.

B. Support for Specialized Services

One of the important capabilities of 5G is the support for specialized services that will rely on connectivity that goes beyond the general network performance or properties of the basic connectivity as offered by the Internet access service. Such support is expected to unleash a whole new range of innovations both at networking and at application service level, innovations for the consumer market or within a special vertical or public sector. However, there are several challenges to the introduction of specialized services, challenges that are, as of today, seemingly in a chicken-and-egg situation.

First, while Europe has now regulations [7] in place addressing the neutrality of the general Internet access service, there are still many uncertainties around what are acceptable implementations of specialized services and how they can be balanced with the Internet access service, in spite of the BEREC guidelines [8]. These uncertainties (among others) have so far kept the network operators reluctant to enter into innovations with new specialized services. While the specialized services can offer better quality and enable richer customer choices, as well as they can unleash innovation by new actors, their realization must be cost-efficient to ensure network operator investments. This implies the need for a balanced approach that can support dynamic and real-time adaptation of resource allocations along the changes in the demands (loads) across the variety of user needs [9].

As a second point, there is a need for both informing the user (or application) about available network services and their offered service levels in real time as well as letting applications signal to the network about their needs in case of value added connectivity (VAC). With the coming 5G era, we foresee a clearer separation between the network service provider (NSP) role and the online application service provider (OAP) role. We anticipate an API offered by the NSPs for letting the OAP invoke and handle VAC for their application service users. In order to ensure that such a VAC API becomes successful, it must be simple but still sufficient in terms of capability support, and standardization is a must. While such an API will establish the common ground, there will be a large space for innovation and competition in how NSPs implement and support such API.

The last challenge to be mentioned here is the need for renewed protocols that can better support the variety of application properties anticipated and explicitly expressed with the specialized services and their VAC API. The ETSI Next Generation Protocols industry specification group (NGP ISG) is investigating scenarios, key performance indicators and new reference models as it identifies current shortcomings and protocol requirements for the future.

C. Multi-access connectivity

5G is made up of several different radio access technologies. While it is customary today for smart personal devices to

support WiFi and cellular, dedicated IoT devices—as part of vehicles, drones, or as stand-alone units—are expected to take this idea even further and support a broader choice of networks, both licensed and unlicensed, further including, e.g., NB-IoT/LPWAN, visible light communication, or mmWave.

The benefits of multi-access connectivity such as increased capacity, robustness or load-balancing have lately been realised with the development of new multipath transport and application protocols (e.g., MPTCP [10], Multipath RTP [11], Multipath QUIC [12]). These mechanisms incorporate signalling means to discover and exchange information about new addresses/paths as well as for load balancing (resource pooling) and failover. In the context of multi-access connectivity, optimizing end-to-end performance requires considering aspects such as best available technology/path selection, best available protocol selection and optimal packet scheduling. However, given the very heterogeneous characteristics of the available access technologies in terms of capacity, delay and loss, optimizing end-to-end performance while providing Quality of Service (QoS) and Quality of Experience (QoE) guarantees remains a challenge. The problem becomes even more challenging especially under mobility scenarios where link characteristics are very dynamic and short-lived.

D. Device to device communication in 5G

D2D is one of the disruptive technologies identified in the roadmap towards 5G since the early stages of discussion [13]. It would allow more efficient use of radio resources, opportunistic access to spectrum, low latency and offer offloading features [14]. Incorporating D2D communication is one new feature of 5G, although it has been longly discussed in Third Generation Partnership Project (3GPP) since release 12, for disaster recovery applications [15], and in release 13, which introduces D2D for relay and public safety applications [16]. Moreover, D2D has been always considered a powerful tool to discover and enable commercial exploitation of local network services. Indeed, the term commonly used to indicate D2D in 3GPP is Proximity Services (ProSe).

D2D will be supported in 5G in terms of both “inband” and “outband” modes, i.e., it will be possible to use the same frequencies used by regular evolved Node B (eNB)-User Equipment (UE) transmissions for UE-UE communications and it will be also possible to establish UE-UE connections by using the eNB as a controller and then switch to an unlicensed channel for actual transmission. The analysis of D2D features in 4G/5G networks is particularly challenging because D2D transmissions are not fully under the control of the cellular operator. Moreover, assessing D2D services requires access to both eNB control messages and to be exposed to D2D transmissions, possibly being involved in D2D exchanges, or being a normal UE suffering interference from D2D links.

E. Internet of Things in 5G

The use of connected devices and applications of smart sensors are ever increasing and will soon influence virtually every aspect of our lives. Out of the over 30 billion connected

devices predicted by 2023 [17], around 20 billion devices are expected to be related to the IoT. 5G is a major driver of this development since enabling IoT and machine-type-communications (MTC) are key 5G use cases.

The requirements of MTC applications are diverse and very different from traditional human-type mobile broadband (MBB) applications. Massive MTC (mMTC) applications like smart cities, asset tracking, or agriculture require scalable and efficient connectivity for a massive number of devices sending infrequent and very short packets. Other applications like self-driving vehicles, smart-grid control, or industrial automation require low latency and very high reliability.

The massive number of devices, as well as the diversity in their requirements, poses challenges for the mobile network infrastructure. The Narrow-band IoT (NB-IoT) technology [18] as well as the LTE-M technology [19] have been proposed by 3GPP in response to some of these challenges. NB-IoT is a massive Low Power Wide Area (LPWA) technology for data perception and acquisition intended for intelligent low-data-rate applications. NB-IoT supports massive connections, ultra-low power consumption, wide area coverage and bidirectional triggering between signaling plane and data plane.

As commercial launches of massive IoT technologies (both NB-IoT and LTE-M) are emerging [20], understanding and characterizing the performance of these technologies from a systems perspective is vital for the successful integration of IoT applications into the 5G ecosystem. Energy-efficiency must be evaluated considering the system as a whole, including all protocol layers and their interaction, and the provided performance in terms of reliability, latency and throughput must be evaluated based on use case traffic patterns and performance requirements.

IV. END-TO-END MEASUREMENTS IN 5G ERA

In this section, we describe how MONROE offers novel opportunities to tackle the above described 5G challenges.

A. Software-Defined Networking and Network Function Virtualization

Clearly the flexibility of 5G, with SDN and NFV that enable slicing and fine grained control of the network, calls for automated mechanism to monitor, control and identify changes, being those due to failures, anomalies, mis-configuration, or eventually malicious attacks. Novel Key Performance Indicators (KPIs) have been defined in the context of 5G, whose definition is a challenge. Indeed, these KPIs represent a shift from the traditional QoS and QoE metrics typically defined and used in the past. 5G KPIs are grouped into three categories: i) Operating KPIs, e.g., “leverage and SME participation”; ii) Performance KPIs, e.g., “1,000 times higher wireless capacity and zero perceived downtime”; and iii) Societal KPIs, e.g., “advanced user controlled privacy and lower energy consumption”. Measuring even the most network oriented KPIs constitutes a challenge due to their high-level definition and heterogeneity. For instance, 49 Performance KPIs have been shortlisted, whose measurability is questionable. Examples

includes “Reducing the average service creation time cycle from 90 hours to 90 minutes (as compared to the equivalent time cycle in 2010)”, “Secure, reliable and dependable Internet with a *zero perceived* downtime for services provision”, “Improved network reliability and resilience”, “End-to-end latency reduced by a factor of 5”, etc. All these need to be rigorously defined, and mapped onto actual measurements that the network can expose.

The experience of MONROE will be key to complete the definition of actually measurable quantities in the context of 5G. Given network management is now cornerstone, measurements are the first step toward visibility and control. SDN, NFV and slicing calls for as much as possible automatic mechanisms based on artificial intelligent approaches to automatically support the network administrators. Verifying Service-Level Agreement (SLA) of dynamic virtual slice is much complicated than the already complex verification of SLA in traditional networks. Accurate and fine grained mechanisms are required to continuously monitor each single element of the network to react to possible problems. Here the MONROE experience in the development of objective methodologies to measure QoS and QoE metrics is fundamental toward the composition of more complicated SLA.

B. Support for Specialized Services

In 5G, there will be a richer set of applications, protocols, and protocol features that need further exploration and evolution. This will in particular be the case for specialized services whether the end-user or end-customer requirements or even SLAs are more explicitly expressed than those of the basic Internet access service. The MONROE platform offers unique capabilities for explicit comparison of network services as well as of protocol features. Moreover, MONROE offers the platform capabilities that will become more important for the public and general communities and their need for transparency with 5G and specialized services. There is a need to ensure proper insights in how NSPs will implement and enforce their commitments to comply with net neutrality regulations along with the clarifications that are anticipated. In order to deal with net neutrality, we anticipate that this will call for explicit commitments by the NSPs on expressing the minimum network performance level in typical busy hours of a set of archetypical applications as achievable by the basic Internet access service. These commitments will naturally vary from a cell or area to another as different areas have different capacities and resources allocated to them. If unexpected events are occurring, the performance might even drop below the expressed target. This will allow the resources for specialized services and for Internet access service to be dynamically allocated according to the load, while the specialized services will also be constrained by session admission control in order to stay within the minimum network performance bounds of the Internet access service.

With these in mind, the MONROE platform will become instrumental in documenting that the NSPs actually do meet their expressed targets, and to what extent they perform

better or in some cases worse than their targets, even for application-specific targets. While the regulators will need such documentations, also public information must be made available as well as NSP specific private information. The MONROE platform is able to handle all these perspectives by careful data storage and privacy handling. Last but not least, the MONROE platform will be able to handle measurement in relation to user specific service level requirements as expressed via the anticipated VAC APIs.

C. Multilink connectivity

The MONROE platform currently provides options of multiple link access technologies. Nodes are connected with multiple LTE modems, and with WiFi and ethernet where available. Thus, using LTE plus WiFi connectivity mimics multi-homing in smart phones with both MBB and WiFi interfaces. MONROE further provides context information (e.g., signal strength, technology, locations, etc...) on the access technology along with measurements. Finally, by building the system over linux and allowing kernel modifications, MONROE enables wide range of protocol experiments. Therefore, MONROE is the ideal platform for experimenting with protocols and applications 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.

To support 5G, the MONROE hardware can easily be extended to support access technologies such as 5G NR, IoT and satellite, integrated using industry standard mPCIe connections, or USB dongles. The MONROE software is portable to other hardware platforms as well, as long as they support standard Linux distributions and is of sufficient performance. In a similar fashion, the software can be extended to provide context information for these technologies. This allows a wide range of experimentation in multilink connectivity in 5G, from enhanced methods to assess the performance of the available networks to improved scheduling algorithms and novel multipath protocols.

D. Device to device communication in 5G

MONROE offers the unique feature of running measurement tests from inside the access network. Therefore, once MONROE nodes will be endowed with 5G NR and fully fledged 5G features, including D2D [14], it will be possible to run both active and passive experiments with 5G D2D. Specifically, MONROE nodes will be able to be engaged in D2D communications, both using inband, under the control of the 5G eNB, and outband, e.g., by using the WiFi interface for direct communications between 5G UEs. This will allow active end-to-end measurements with D2D users in 5G cells. As concerns passive experiments, MONROE nodes will be able to observe D2D control traffic coming from the eNB and partially observe direct transmissions from D2D nodes. Therefore, the MONROE platform will offer the possibility to evaluate the quality of ProSe services and relay services while, at the same time, offering the possibility of assessing the

impact of D2D transmissions on cellular users (i.e., on speed, efficiency, costs, latency, etc.) in an operational 5G system.

E. Internet of Things in 5G

The MONROE platform targets measurements that capture performance from an end-device and application perspective, allowing both long-term monitoring of KPIs and select use case experiments. Once extended with support for massive IoT technologies, MONROE will thus offer possibilities for long-term monitoring of IoT network performance as well as targeted performance studies. The experiment as a service (EaaS) toolset developed for 4G networks within MONROE can be extended to cover important 5G IoT use cases, allowing to monitor and assess the IoT network performance against IoT use case requirements. Such studies are highly important as while the basic network-oriented KPIs for 5G are defined, their validation in operational systems and their impact on end-user performance remains open.

MONROE measurements will capture the performance of both the IoT protocol stack and the 4G/5G protocol stack, as well as their interactions, allowing to assess energy efficiency and performance of the system as a whole. This is important not only for use case experiments as discussed above, but also allows detailed studies of protocol performance and interactions. It will for instance allow targeted studies of how to optimize the interaction between IoT protocols such as Constrained Application Protocol (CoAP) [21] and the underlying 4G/5G protocol stack for improved energy efficiency, where the settings of different timers and other protocol parameters can have a large impact on the energy consumed. Such interactions have so far only been studied analytically, making evaluation and optimization in operational conditions key for predicting the lifetime of IoT devices and reaching desired energy efficiency targets.

V. CONCLUSION

End-to-end network measurements are essential resources in many network investigations, especially for performance and reliability analysis of complex networks and, applications and services that are running on these networks, since they provide an environment that is hard to mimic in models and simulators. Supporting a wide range of 5G use cases demands high speeds, ultra-low latency, ultra-high reliability, massiveness of network scenarios. MONROE has proven to be a key EaaS platform for validation of 4G KPIs. In this paper, we argue that such a platform and a set of common methodologies will be even more important and more challenging with 5G and we describe the opportunities a 5G-capable MONROE platform can provide to address the key challenges of 5G in the context of end-to-end measurements.

ACKNOWLEDGMENT

This work is funded by the EU H2020 research and innovation programme under grant agreement No. 644399 (MONROE), and by the Norwegian Research Council project No. 250679 (MEMBRANE). The work of V. Mancuso was

supported by a Ramon y Cajal grant (ref: RYC-2014-16285) from the Spanish Ministry of Economy and Competitiveness.

REFERENCES

- [1] Cisco, "Cisco Visual Networking Index: Forecast and Methodology, 2016–2021," CISCO, Tech. Rep., September 2017.
- [2] ITU-R, "IMT Vision - Framework and overall objectives of the future development of IMT for 2020 and beyond," Sep. 2015.
- [3] 5G PPP Architecture Working Group and others, "View on 5G architecture," *White Paper*, Jul. 2016.
- [4] Ö. Alay, A. Lutu, M. Peón-Quirós, V. Mancuso, T. Hirsch, K. Evensen, A. Hansen, S. Alfredsson, J. Karlsson, A. Brunström, A. Safari Khatouni, M. Mellia, and M. Ajmone Marsan, "Experience: An Open Platform for Experimentation with Commercial Mobile Broadband Networks," in *Proc. of ACM Mobicom.*, Oct. 2017.
- [5] C. Midoglu, L. Wimmer, A. Lutu, O. Alay, and C. Griwodz, "MONROE-Nettest: A Configurable Tool for Dissecting Speed Measurements in Mobile Broadband Networks," in *Proc. of IEEE CNERT (INFOCOM Workshop)*, Apr. 2018.
- [6] A. Finamore, M. Mellia, M. Meo, M. M. Munafo, P. D. Torino, and D. Rossi, "Experiences of internet traffic monitoring with tstat," *IEEE Network*, vol. 25, no. 3, pp. 8–14, May 2011.
- [7] European Union, "Regulation (EU) 2015/2120 of the European Parliament and of the Council of 25 November 2015 laying down measures concerning open internet access and amending Directive 2002/22/EC on universal service and users rights relating to electronic communications networks and services," *Official Journal of the European Union*, vol. L 310, pp. 1–18, 2015-11-26. [Online]. Available: <http://data.europa.eu/eli/reg/2015/2120/oj>
- [8] Body of European Regulators for Electronic Communications, "BEREC Guidelines on the Implementation by National Regulators of European Net Neutrality Rules," Aug. 2016. [Online]. Available: http://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/6160-berec-guidelines-on-the-implementation-b_0.pdf
- [9] H. Lønsethagen, (Eds) *et al.*, "Service Level Awareness and open multi-service internetworking," 2016. [Online]. Available: <https://www.networld2020.eu/sria-and-whitepapers/>
- [10] D. Wischik, C. Raiciu, A. Greenhalgh, and M. Handley, "Design, Implementation and Evaluation of Congestion Control for Multipath TCP," in *Proc. of USENIX NSDI*, Mar. 2011.
- [11] E. Pakulova, K. Miller, and A. Wolisz, "Adaptive low-delay video streaming in heterogeneous wireless networks using MPRTp," in *Proc. of IWCMC*, Jun. 2017.
- [12] T. Viernickel, A. Frömmgen, A. Rizk, B. Koldehofe, and R. Steimetz, "Multipath QUIC: A Deployable Multipath Transport Protocol," in *Proc. of IEEE ICC*, 2018.
- [13] F. Boccardi, R. W. Heath, A. Lozano, T. L. Marzetta, and P. Popovski, "Five disruptive technology directions for 5G," *IEEE Communications Magazine*, vol. 52, no. 2, pp. 74–80, Feb. 2014.
- [14] A. Asadi, Q. Wang, and V. Mancuso, "A Survey on Device-to-Device Communication in Cellular Networks," *IEEE Communications Surveys Tutorials*, vol. 16, no. 4, pp. 1801–1819, Fourth quarter 2014.
- [15] 3GPP TS 23.303, Release 12, "Technical Specification Group and System Aspects, Proximity-based services (ProSe); Stage 2," Dec. 2014.
- [16] 3GPP TS 23.303, Release 13, "Proximity-Based Services (ProSe); Stage 2," Jun. 2015.
- [17] P. Jonsson, S. Carson, J. S. Sethi, M. Arvedson, R. Svenningsson, P. Lindberg, K. Öhman, and P. Hedlund, "Ericsson Mobility Report: November 2017," Ericsson, Technology and Emerging Business, Tech. Rep., Nov. 2017.
- [18] Y. P. E. Wang, X. Lin, A. Adhikary, A. Grovlen, Y. Sui, Y. Blankenship, J. Bergman, and H. S. Razaghi, "A Primer on 3GPP Narrowband Internet of Things," *IEEE Communications Magazine*, vol. 55, no. 3, pp. 117–123, Mar. 2017.
- [19] M. Elsaadany and W. Hamouda, "The new enhancements in LTE-A Rel-13 for reliable machine type communications," in *Proc. of IEEE PIMRC*, Oct. 2017.
- [20] GSM Association, "Mobile IoT Rollout Report," 2018. [Online]. Available: <https://www.gsma.com/iot/miot-rollout/>
- [21] Z. Shelby, K. Hartke, and C. Bormann, "The Constrained Application Protocol (CoAP)," Internet Requests for Comments, RFC 7252, June 2014. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc7252.txt>