

Exploring QoE in cellular networks: how much bandwidth do you need for popular smartphone apps?

Pedro Casas, Raimund Schatz, Florian Wamser, Michael Seufert, Ralf Irmer

Angaben zur Veröffentlichung / Publication details:

Casas, Pedro, Raimund Schatz, Florian Wamser, Michael Seufert, and Ralf Irmer. 2015. "Exploring QoE in cellular networks: how much bandwidth do you need for popular smartphone apps?" In Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges, 17 August 2015, London, UK, edited by David Choffnes and Kobus Van der Merwe, 13-18. New York, NY: ACM. <https://doi.org/10.1145/2785971.2785978>.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



Exploring QoE in Cellular Networks: How Much Bandwidth do you Need for Popular Smartphone Apps?

Pedro Casas (1), Raimund Schatz (1), Florian Wamser (2), Michael Seufert (2), Ralf Irmer (3)
(1) FTW Vienna, (2) University of Würzburg, (3) Vodafone Research and Development
{casas|schatz}@ftw.at; {florian.wamser|seufert}@informatik.uni-wuerzburg.de; ralf.irmer@vodafone.com

ABSTRACT

A quarter of the world population will be using smartphones to access the Internet in the near future. In this context, understanding the Quality of Experience (QoE) of popular services in such devices becomes paramount for cellular network operators, who need to offer high quality levels to reduce the risks of customers churning for quality dissatisfaction. In this paper we study the problem of QoE provisioning in smartphones, presenting the results obtained from subjective lab tests performed for five popular apps: YouTube, Facebook, Web browsing through Chrome, Google Maps, and WhatsApp. The analysis addresses the impact of the access downlink bandwidth on the QoE of these apps when accessed through smartphones. The results presented in this paper provide a sound basis for better understanding the QoE requirements of popular services and mobile apps, as well as for dimensioning the underlying provisioning network. To the best of our knowledge, this is the first paper providing such a comprehensive analysis of QoE in mobile devices.

Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: Miscellaneous;
C.4 [Performance of Systems]: Measurement Techniques

Keywords

QoE; Smartphones; Subjective Lab Tests; Mobile Apps

1. INTRODUCTION

Smartphones are becoming the most typical mobile device to access Internet today. Recent projections [2] show that by 2016, a quarter of the world population will be using smartphones to access the most popular services such as YouTube, Facebook and WhatsApp. According to Cisco's global mobile data traffic forecast [1], smartphones will be responsible for more than three-quarters of the mobile data traffic generated by 2019. In the light of these trends, cellular network

operators are becoming more and more interested in understanding how to dimension their access networks and how to manage their customers' traffic to capture as many new customers as possible. In this scenario, the concept of Quality of Experience (QoE) has the potential to become one of the main guiding paradigms for managing quality in cellular networks. Closely linked to the subjective perception of the end-user, QoE enables a broader, more holistic understanding of the factors that influence the performance of systems, complementing traditional technology-centric concepts such as Quality of Service (QoS). Indeed, QoE is today an important differentiator between providers, but most of the times, operators do not really grasp the key aspects related to QoE in their networks.

In this paper we claim that understanding QoE in mobile devices is paramount for cellular network operators, and present the results obtained from subjective lab tests performed for popular end-user services accessed through smartphones. In particular, we consider the following five well-known applications in mobile devices: YouTube dynamic and non-dynamic video streaming, Facebook, Google Maps (Gmaps from now on), Web browsing through Google Chrome, and WhatsApp. The evaluations performed in these subjective tests consider the impact of the most relevant QoS-based characteristics of the access network: the downlink bandwidth. The main contribution of our study is to shed light into the problem of QoE-based network provisioning for mobile devices, offering a comprehensive analysis of the QoE undergone by users when the underlying access network presents different QoS characteristics or performance levels.

The standard approach to assess the performance of networks and services from a QoE end-user perspective is to conduct controlled lab experiments [19–21]. The key benefits of such an approach rely on the full control the experimenter has on the overall evaluation process. Indeed, content and context are fully known and controlled, and users are directly briefed and observed on the spot, providing as such tangible and solid results.

The remainder of the paper is organized as follows: Sec. 2 presents an overview of the related work on QoE for web and cloud-based services, focusing on the specific case of mobile devices. Sec. 3 describes the experimental setup used in the QoE subjective tests. Sec. 4 presents the main results of the study, including the impact of the downlink bandwidth on the level of satisfaction and overall experience of the end-user when accessing the aforementioned applications through a smartphone. Sec. 5 discusses and concludes this work.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in:

AllThingsCellular'15, August 17-21, 2015, London, United Kingdom

© 2015 ACM. ISBN 978-1-4503-3538-6

DOI: <http://dx.doi.org/10.1145/2785971.2785978>

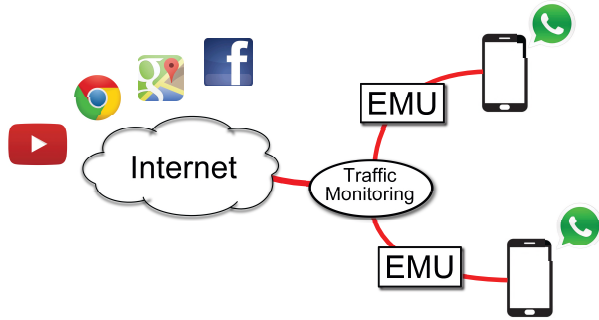


Figure 1: Experimental setup used in the study. Devices are connected to the Internet through independent, controlled WiFi connections.

2. RELATED WORK

The study of the QoE requirements for web-based services and cloud-based applications as the ones we target in this paper has a long list of fresh and recent references. A good survey of the QoE-based performance of mobile networks when accessing many different web and cloud services is presented in [7]. The main limitation of previous work when considering the analysis of performance in cellular networks is that the considered access devices are not smartphones, but rather traditional laptops with mobile broadband connections.

The specific case of QoE in YouTube deserves particular attention, due to the overwhelming popularity and omnipresence of the service. Studies have both considered the “standard” HTTP video streaming flavour of YouTube, as well as the more recent Dynamic Adaptive Streaming (DASH) version. Previous papers [9, 10] have shown that stalling (i.e., stops of the video playback) and initial delays on the video playback are the most relevant Key Performance Indicators (KPIs) for QoE in standard HTTP video streaming. In the case of adaptive streaming, a new KPI becomes relevant in terms of QoE: quality switches. In particular, authors in [12] have shown that quality switches have an important impact on QoE, as they increase or decrease the video quality during the playback. A comprehensive survey of the QoE of adaptive streaming can be found in [13].

When it comes to our specific analysis of QoE in mobile networks and mobile devices, references become scarcer, showing that there is still an important gap to fill. In [15], authors study the characteristics of YouTube traffic on smartphones connected to a cellular network, showing that these devices have a non-negligible impact on the characteristics of the downloaded traffic. Closer to the subject of this paper, authors in [16] describe a subjective QoE evaluation framework for mobile Android devices in a lab environment. In [17], authors study the QoE of YouTube in mobile devices through a field trial, exclusively considering the non-adaptive version of the YouTube player. Authors in [18] recently introduced Prometheus, an approach to estimate QoE of mobile apps, using both passive in-network measurements and in-device measurements, applying machine learning techniques to obtain mappings between QoS and QoE. Additional papers in a similar direction tackle the problem of modeling QoE for Web [4] and video [5].

Table 1: Operational expected RTT and downlink bandwidth values for different access technologies.

| Access Technology | RTT (ms) | Downlink Bandwidth |
|-------------------|----------|--------------------|
| LTE | < 50 | 20 Mbps |
| HSPA+ | < 50 | 10 Mbps |
| HSPA | < 150 | 4 Mbps |
| UMTS | < 200 | 384 kbps |
| EDGE | < 350 | 160 kbps |
| GPRS | < 650 | 40 kbps |

Finally, WhatsApp is a very new service and its study has been so far quite limited. Some recent papers have partially addressed the characterization of its traffic, including a QoE perspective [14].

3. EXPERIMENTAL METHODOLOGY

The subjective study consists of 52 participants interacting with the aforementioned services while experiencing different downlink bandwidth profiles in the background data connection. Figure 1 depicts a high-level diagram of the experimental testbed employed in the subjective tests. Android smartphone devices are used in the study (Samsung Galaxy S4, OS Android 4.4 KitKat). Devices are connected to the Internet through separate WiFi access networks. The downlink traffic between the different evaluated services and the devices is routed through a modified version of the very well known NetEm network emulator [23] so as to control the different access network profiles under evaluation.

Different bandwidth profiles are instantiated at the network emulators, changing downlink bandwidth from 0.5 Mbps to 16 Mbps. These profiles are selected from operational experience, particularly following typical operational values reported in Table 1 for different access network technologies (LTE, 3G/2G, etc.). The list includes both research results (e.g., [6]) as well as operational knowledge coming from cellular operators, collaborating with the project which drives this work, the ACE project¹. Access RTT is kept at 10 ms when downlink bandwidth is varied, which corresponds to optimal performance in LTE and evolved networks.

Participants were instructed to perform independent tasks for each of the five considered applications. For YouTube, they were requested to watch two-minutes HD YouTube videos, considering both the usage of the standard (i.e., non-DASH) and the DASH versions of the YouTube player. Table 2 reports the YouTube videos used (YouTube video IDs), which include four mainly nature-themed clips and a movie trailer. Videos correspond to 4K ultra-HD videos (i.e., 2160p), which are down-scaled to HD resolution (i.e., 720p) due to the device’s display capabilities (i.e., screen size and resolution). To better understand the obtained results, the table also reports the average video bit rate of the corresponding videos, which is in all cases around 1.6 Mbps. In the case of Facebook, participants were instructed to access the application with a specific user account, browse the timeline of this user, and browse through specific photo albums created for this user. In the WhatsApp tests, participants worked in couples and exchanged specific video files of fixed size (i.e., 5 MB), and the participant downloading the video

¹The ACE project - FTW Vienna, <http://ace.ftw.at/>

Table 2: Video content and bitrate.

| YouTube Video ID | Average Video bitrate |
|------------------|-----------------------|
| 6pxRHBw-k8M | 1.5 Mbps |
| iNJdPyoqt8U | 1.5 Mbps |
| kObNpTFPV5c | 1.7 Mbps |
| QS7lN7giXXc | 1.7 Mbps |
| suWsd372pQE | 1.6 Mbps |

file was the one providing a QoE evaluation, based on the experienced time. Web browsing tasks consisted of reading and browsing through a popular and complex News website (<http://edition.cnn.com/>). Finally, Gmaps tasks consisted of exploring different city maps using the Gmaps application, in satellite view, which consumes more bandwidth.

Tests were performed in a dedicated lab for subjective studies, compliant with the QoE subjective studies standards [19–21]. All traffic flows are captured and exported to standard pcap traces for off-line traffic analysis, using high-performance Endance DAG cards. Regarding participants’ demographics, 29 participants were female and 23 male, the average age was 32 years old, with 40 participants being less than 30 years old. Around half of the participants were students and almost 43% were employees, and 70% of the participants have completed university or baccalaureate studies.

Regarding QoE feedback, participants were instructed to rate their *overall experience* (rate the overall quality) according to a continuous ACR Mean Opinion Score (MOS) scale [19], ranging from “bad” (i.e., MOS = 1) to “excellent” (i.e., MOS = 5). MOS ratings were issued by participants through a custom questionnaire application running on separate laptops, which pops up immediately after a condition was tested. Participants also provided feedback on the *acceptability* of the application, stating whether they would continue using the application under the corresponding conditions or not. For the specific case of YouTube, three additional questions were asked to participants: (i) *initial playback delay annoyance* (did you perceive the initial loading time of the video as disturbing?); (ii) *stalling annoyance* (did you perceive stalling as disturbing?); (iii) *video image quality* (rate the image quality of the video). Each testing session runs for a total time of two hours. Participants were compensated with vouchers for their participation, which proved to be sufficient for achieving correct involvement in the tasks.

4. QOE IN MOBILE DEVICES

In this section we present and discuss the results obtained in the conducted tests. Constant downlink bandwidth profiles are tested for the five studied services. In the case of YouTube, the Downlink BandWidth (DBW) takes the values 1 Mbps, 2 Mbps, and 4 Mbps (recall that the average video bitrate of the tested videos is around 1.6 Mbps). Facebook is tested with DBW = 0.5 Mbps, 1 Mbps, 2 Mbps, and 8 Mbps. The profiles for Web browsing are almost identical to those used in Facebook, expect that the last condition corresponds to an optimal DBW = 16 Mbps. Gmaps is tested with a fully logarithmic scale: 1 Mbps, 2 Mbps,

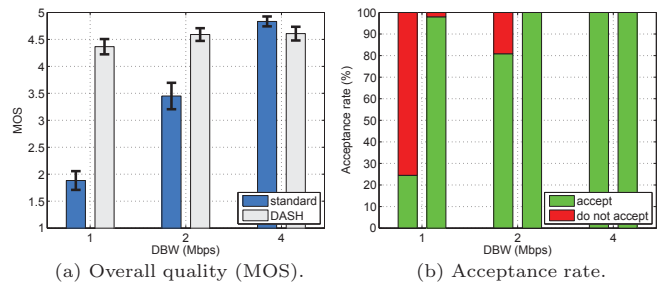


Figure 2: QoE in YouTube standard (i.e., non-DASH) and DASH. Overall quality and acceptability for different downlink bandwidth configurations.

4 Mbps, 8 Mbps, and 16 Mbps. Finally, the WhatsApp DBW profile takes the values 0.5 Mbps, 1 Mbps, 2 Mbps, 4 Mbps, and jumps to 16 Mbps to verify the occurrence of QoE-saturation, which we shall explain next.

A final remark regarding interpretation of results: the reader shall note that the maximum MOS ratings declared by the participants are never 5 but somewhere between 4.2 and 4.6. This is a well known phenomenon in QoE studies called *rating scale saturation*, where users hardly employ the limit values of the scale for their ratings [7]. So from now on, we shall consider as optimal quality a MOS score close to 4.5.

4.1 QoE in YouTube Mobile

Figure 2 reports the overall quality and acceptability results obtained for the YouTube tests. Recall that in the YouTube scenario, we compare the standard, non-adaptive version of the YouTube player (videos are selected to play in HD quality) against the DASH-capable one. In the DASH case, videos are also requested in HD quality, but the server adapts the subsequent video quality resolutions to the bandwidth estimated by the player.

Figure 2(a) compares the overall QoE experienced by the participants using both player versions. It is quite impressive to appreciate how the DASH approach results in a nearly optimal QoE for all the tested conditions (from 1 Mbps to 4 Mbps), whereas the fixed HD quality approach results in poor QoE for downlink bandwidth below 4 Mbps. As expected for the standard player, heavy stalling occurs for the 1 Mbps condition, taking into account that the average video bitrate is 1.6 Mbps. Indeed, as we have shown in [22], the DBW should be in the order of 30% higher than the average video bitrate to avoid stalling when non-adaptive streaming is used. This dimensioning rule also explains the results obtained for the 2 Mbps condition, as some stalling still occurs. No stalling seems to occur for the DASH version. The main difference is that DASH changes the video quality without incurring in playback stalling, whereas the fixed quality configuration definitely results in video stalling.

Figure 2(b) reports the results in terms of acceptability of the participants. This is one of the key features that an operator has to consider, because low acceptance rate may sooner or later turn into churn. As observed, acceptance rate is as low as 23% for the standard streaming at 1 Mbps, whereas it’s close to 99% in the case of DASH.

To complement the picture for YouTube QoE in mobile devices, Figure 3 depicts the results obtained in terms of (a)

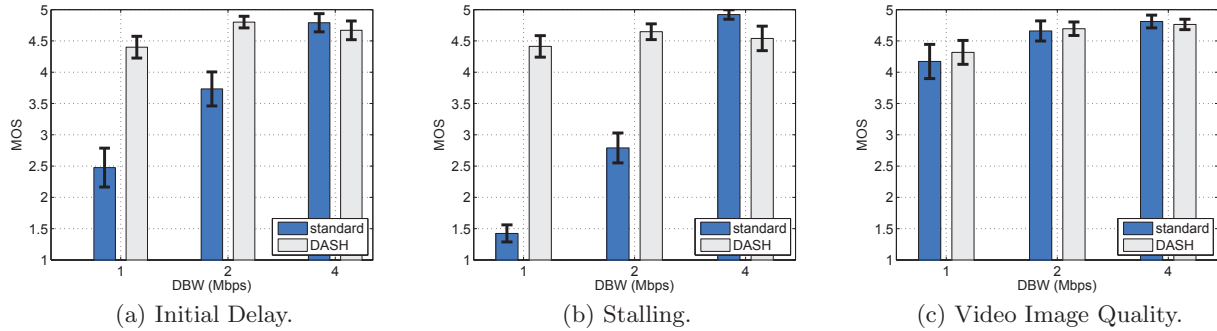


Figure 3: QoE for YouTube Mobile, considering DASH and non-DASH. Videos are UHD 4k, but due to the device capabilities they are re-scaled to 720p, resulting in an average video bit rate of around 1.6 Mbps for all the considered videos.




Figure 4: QoE in Facebook. Overall quality and acceptability for different DBW




Figure 5: QoE in Web browsing (news website). Overall quality and acceptability for different DBW.

annoyance caused by the initial delay (time to first frame displayed), (b) annoyance caused by stalling (stop of the video playback), and (c) video image quality. In Figs. 3(a) and 3(b), a MOS = 5 means not disturbing at all, whereas a MOS = 1 means unbearable (very annoying). Both the initial delay and the stalling impact on QoE follow the same pattern as the overall QoE for the non-adaptive application. However, stalling has a much stronger impact on the user's level of annoyance, confirming what has been already seen in previous studies for desktop and laptop like devices. It is very interesting the fact that the usage of DASH also reduces significantly the initial delay, suggesting that the quality of the video chunks assigned by the server and the player is estimated from active measurements prior to the playback, rather than on top of the video traffic itself.

The most interesting and remarkable results is presented in Figure 3(c), which reports the perceived image quality of the video. According to previous studies [12], quality switches induced by DASH have an important impact on QoE. However, in the case of smartphones, where displays are smaller than laptops or desktop devices, quality switches do not seem to have an important impact on the perception of the user. While these results are directly linked to the specific quality-switching patterns induced by the tested DBW conditions, they show some potentially remarkable contribution to assess QoE for YouTube in smartphones when using DASH. As a summary, using DASH reduces both the chances of stalling and the initial video playback delays, at no apparent perceived image quality cost.

4.2 QoE in Facebook Mobile

Figure 4 reports the results obtained in the Facebook tests for different DBW configurations, considering both (a) the overall quality and (b) the acceptance rate. A DBW of 500 kbps is not high enough to reach full user satisfaction in Facebook mobile for Android devices, as participants declared a fair quality with an acceptance rate of about 80%. Still, a DBW of 1 Mbps results in good overall quality, with almost full acceptance of the participants. Excellent QoE results are attained for 8 Mbps, which shows that even if a 2 Mbps DBW allocation is high enough to reach full acceptance (cf. Figure 4), the overall experience of the user can still marginally improve. These DBW thresholds explain the boundaries between user satisfaction and resources over-provisioning. Interestingly, these QoE DBW requirements are more restrictive than those we found in [3] for laptops about 2 years ago, evidencing how the Facebook app is becoming more network resources demanding.

4.3 QoE in Mobile Web Browsing

Figure 5 reports the overall quality and acceptability results obtained for the News website browsing tests. Note first how the quality increases in a logarithmic fashion with increasing values of the DBW. Good experience (MOS \approx 4) is obtained for a DBW of 2 Mbps, and only slight QoE differences are obtained when increasing the bandwidth to up to 16 Mbps, going to MOS \approx 4.15. Going in the DBW decreasing direction, the slowest tested condition still results in fair quality (MOS \approx 3.5) and high acceptance rate, close to 90%.




Figure 6: QoE in Gmaps. Overall quality and acceptability for different DBW.

4.4 QoE in Gmaps Mobile

Figure 6 reports the overall quality and acceptability results obtained for the Gmaps tests. Figure 6(a) shows that a DBW of 4 Mbps results in near optimal QoE (MOS \approx 4.5), and from this value on, QoE saturation already occurs. This means that no major QoE improvements are then obtained for additional bandwidth provisioning. A DBW of 2 Mbps provides good quality results and almost full acceptance, but a DBW of 1 Mbps rapidly brings Gmaps into bad user experience.

4.5 QoE in WhatsApp

Figure 7 shows the QoE results for different DBW values. Users tolerate WhatsApp downloads with a good overall experience and high acceptability as long as the DBW is above 2 Mbps, but user experience heavily degrades for slower connections, resulting in very bad quality for a DBW of 500 kbps. In this case, a DBW threshold of 2 Mbps permits to approximately discriminate between good and bad experience. Given the file size used in the tests (5 MB), there is a clear saturation effect after 4 Mbps, as QoE does not increase for higher DBW values. Finally, even if the obtained results are partially biased by both the specific file size used in the tests and the participants task briefing, obtained results are similar to those we obtained in [8] for the specific case of Dropbox file sharing, suggesting that the main take aways are potentially more generic than expected when considering file downloads, either in mobile devices or in fixed ones.

5. CONCLUDING REMARKS

Smartphones are becoming the Internet-access devices by default, and we claim that network operators must understand how to manage and dimension their networks in order to correctly provision popular services accessed in smartphones, avoiding wasting additional unnecessary resources while keeping end users happy, and most importantly, reducing the chances of churning due to quality dissatisfaction.

We have presented an overview on the QoE of different services and applications with different network-level QoS requirements for the specific case of smartphone devices. Our results are highly relevant to future 5G design and LTE evolution in better understanding the mapping between network performance and customer experience. Indeed, they are very practical and have a direct impact on the operation and management of mobile networks.

Obtained results suggest that a downlink bandwidth of 4 Mbps is high enough to reach near optimal results in terms of




Figure 7: QoE in WhatsApp. Overall quality and acceptability for different DBW.

overall quality and acceptability for the evaluated services when accessed in smartphones. As a consequence, cellular network operators should target such a downlink bandwidth as their short term goal for dimensioning their access networks. Given this relatively low requirement, resources could be re-allocated or scheduled to manage the network more easily and with a more efficient cost-benefit trade-off, avoiding over-provisioning while keeping high QoE. The implications for the end-user are straightforward: you do not need a super high speed cellular contract with your operator if your target is on the studied applications. So in particular, an expensive LTE contract is not necessary to have a near optimal experience today.

We have also shown that dynamic applications such as YouTube DASH are much better suited to smartphone scenarios, providing the same level of experience as the non-adaptive version of the YouTube application in terms of image quality, but with much lower QoS-based requirements in terms of downlink bandwidth. This is a major finding, as DASH has been shown to degrade the video image quality and the associated user experience when considering standard, laptop or PC devices. The main difference with smartphones is their inherent small size displays, which to some extent filter out the impact of quality switches. A direct implication of this finding is that cellular network operators willing to monitor the QoE of its YouTube customers must know which type of technology is used by the YouTube app in the smartphone to understand the QoE from downlink bandwidth measurements. Indeed, the QoE could be either excellent or very bad for the same measured average downlink bandwidth of 1 Mbps, depending if adaptive or non-adaptive technology is in place.

Finally, we are very aware that our results only tackle one side of the problem: the experience of the customers, from a very simple perspective: the downlink bandwidth. We agree with other researchers in that a more holistic perspective incorporating QoE, energy-consumption, data (re)transmission, and radio resource impact (among others) should be considered. This paper provides some initial components of such a holistic analysis.

Acknowledgment

The research leading to these results has received partial funding from the European Union under the FP7 Grant Agreement n. 318627, "mPlane", and has been done within the project ACE 3.0 at the Telecommunications Research Center Vienna (FTW).

6. REFERENCES

- [1] Cisco Visual Networking Index, “Global Mobile Data Traffic Forecast Update 2014-2019 White Paper”, 2014.
- [2] eMarketer Newsletter, “2 Billion Consumers Worldwide to Get Smart(phones) by 2016”, 2014.
- [3] P. Casas, A. Sackl, S. Egger, R. Schatz, “YouTube & Facebook Quality of Experience in Mobile Broadband Networks”, in *Proc. IEEE Globecom Workshops*, 2012.
- [4] A. Balachandran, A. Vaneet, E. Halepovic, J. Pang, S. Seshan, S. Venkataraman, H. Yan, “Modeling Web Quality of Experience on Cellular Networks”, in *Proc. ACM MOBICOM*, 2014.
- [5] M. Shafiq, J. Erman, L. Ji, A. Liu, J. Pang, J. Wang, “Understanding the Impact of Network Dynamics on Mobile Video User Engagement”, in *Proc. ACM SIGMETRICS*, 2014.
- [6] M. Laner, P. Svoboda, P. Romirer-Maierhofer, N. Nikaen, F. Ricciato, M. Rupp, “A Comparison Between One-Way Delays in Operating HSPA and LTE Networks”, in *Proc. WiOpt*, 2012.
- [7] P. Casas, R. Schatz, “Quality of Experience in Cloud Services: Survey and Measurements”, in *Computer Networks*, vol. 68, pp. 149-165, 2014.
- [8] P. Casas, R. Fischer, S. Suetete, R. Schatz, “A First Look at Quality of Experience in Personal Cloud Storage Services”, in *Proc. IEEE ICC Workshops*, 2013.
- [9] T. Hoffeld, M. Seufert, M. Hirth, T. Zinner, P. Tran-Gia, R. Schatz, “Quantification of YouTube QoE via Crowdsourcing”, in *Proc. IEEE ISM*, 2011.
- [10] R. Mok, E. Chan, X. Luo, R. Chang, “Inferring the QoE of HTTP Video Streaming from User-Viewing Activities”, in *Proc. ACM SIGCOMM W-MUST*, 2011.
- [11] T. Hoffeld, S. Egger, R. Schatz, M. Fiedler, K. Masuch, C. Lorentzen, “Initial Delay vs. Interruptions: Between the Devil and the Deep Blue Sea”, in *Proc. QoMEX*, 2012.
- [12] B. Lewcio, B. Belmudez, A. Mehmood, M. Wältermann, S. Möller, “Video Quality in Next Generation Mobile Networks – Perception of Time-varying Transmission”, in *Proc. IEEE CQR*, 2011.
- [13] M. Seufert, S. Egger, M. Slanina, T. Zinner, T. Hoffeld, P. Tran-Gia, “A Survey on Quality of Experience of HTTP Adaptive Streaming”, in *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 469-492, 2015.
- [14] P. Fiadino, M. Schiavone, P. Casas, “Vivisecting WhatsApp in Cellular Networks: Servers, Flows, and Quality of Experience”, in *Proc. TMA*, 2015.
- [15] J. J. Ramos-Muñoz, J. Prados-Garzon, P. Ameigeiras, J. Navarro-Ortiz, J. M. Lopez-Soler, “Characteristics of Mobile YouTube Traffic”, in *IEEE Wireless Communications*, vol. 21, no. 1, pp. 18-25, 2014.
- [16] I. Ketykó, K. De Moor, T. De Pessemier, A. Verdejo, K. Vanhecke, W. Joseph, L. Martens, L. De Marez, “QoE Measurement of Mobile YouTube Video Streaming”, in *Proc. ACM MoViD*, 2010.
- [17] G. Gómez, L. Hortigüela, Q. Pérez, J. Lorca, R. García, M. Aguayo-Torres, “YouTube QoE Evaluation Tool for Android Wireless Terminals,” in *EURASIP Journal on Wireless Communications and Networking*, vol. 164, 2014.
- [18] V. Aggarwal, E. Halepovic, J. Pang, S. Venkataraman, H. Yan, “Prometheus: Toward Quality-of-Experience Estimation for Mobile Apps from Passive Network Measurements,” in *Proc. ACM HotMobile*, 2014.
- [19] Int. Telecommunication Union, “Methods for Subjective Determination of Transmission Quality”, *ITU-T Rec. P.800*, 1996.
- [20] Int. Telecommunication Union, “Subjective Video Quality Assessment Methods for Multimedia Applications”, *ITU-T Rec. P.910*, 2008.
- [21] Int. Telecommunication Union, “Subjective Testing Methodology for Web Browsing”, *ITU-T Rec. P.1501*, 2013.
- [22] P. Casas, M. Seufert, R. Schatz, “YOUQMON: A System for On-line Monitoring of YouTube QoE in Operational 3G Networks,” in *ACM SIGMETRICS PER*, vol. 41, 2013.
- [23] S. Hemminger, “Network Emulation with NetEm”, in *Proc. 6th Australia’s National Linux Conference – LCA*, 2005.