

Workshop on Tracking Quality of Experience in the Internet: Summary and Outcomes

Fabián E. Bustamante^{*}, David Clark[†] and Nick Feamster[‡],
Northwestern University^{*} MIT[†] Princeton University[‡]

June 22, 2016

Abstract

This is a report on the Workshop on *Tracking Quality of Experience in the Internet*, held at Princeton, October 21–22, 2015, jointly sponsored by the National Science Foundation and the Federal Communication Commission. The term Quality of Experience (QoE) describes a user’s subjective assessment of their experience when using a particular application. In the past, network engineers have typically focused on Quality of Service (QoS): performance metrics such as throughput, delay and jitter, packet loss, and the like. Yet, performance as measured by QoS parameters only matters if it affects the experience of users, as they attempt to use a particular application. Ultimately, the user’s experience is determined by QoE impairments (*e.g.*, rebuffering). Although QoE and QoS are related—for example, a video rebuffering event may be caused by high packet-loss rate—QoE metrics ultimately affect a user’s experience.

Identifying the causes of QoE impairments is complex, since the impairments may arise in one or another region of the network, in the home network, on the user’s device, in servers that are part of the application, or in supporting services such as the DNS. Additionally, metrics for QoE continue to evolve, as do the methods for relating QoE impairments to underlying causes that could be measurable using standard network measurement techniques. Finally, as the capabilities of the underlying network infrastructure continues to evolve, researchers should also consider how to design infrastructure and tools can best support measurements that can better identify the locations and causes of QoE impairments.

The workshop’s aim was to understand the current state of QoE research and to contemplate a community agenda to integrate ongoing threads of QoE research into a collaboration. This summary report describes the topics discussed and summarize the key points of the discussion. Materials related to the workshop are available at <http://aqualab.cs.northwestern.edu/conference/276#program>.

1 Introduction

Quality of Experience (QoE) describes a user’s subjective assessment of the experience when using a particular application or service. Network engineers have in the past focused on Quality of Service (QoS), which entails performance metrics such as throughput, delay, jitter, and packet loss. Yet, conventional QoS metrics are ultimately most consequential when they affect the experience of end users, for a particular application that they are using (*e.g.*, streaming video). A number of research groups have started to investigate better methods, tools, and network infrastructure to measure QoE in today’s Internet, to identify the causes of QoE impairments, and to determine ways to improve QoE. This workshop aimed to understand the state of QoE research and to contemplate a community agenda towards common research goals. Towards this end, the FCC posed a number of questions, including:

- How to define and measure QoE.
- How to relate QoE to measurable network performance metrics (*e.g.*, packet loss, jitter, latency).
- Developing strategies for consistently and continually measuring QoE, from dedicated infrastructure to crowdsourcing.
- Measuring QoE in the presence of constraints, such as limited resources and encryption.

To study these and other questions, we organized the workshop into sessions in a top-down manner, from user measurement studies to network measurement infrastructure. Each session comprised a sequence of talks, followed by open discussion.

We trust that the exchange of ideas and independent research agendas will ultimately yield multiple collaborations aimed at measuring and analyzing the state of QoE in the current Internet and identifying and locating the causes of QoE impairments. This was indeed one of the key goal of the workshop. Locating QoE impairments is a complex task, as these impairments may arise in various regions of the network (ranging from the home to the wide area), on the user’s device, on servers or within data centers, or due to supporting services (*e.g.*, DNS). Ultimately, identifying, locating, and eradicating QoE impairments

will involve collaboration of researchers across multiple areas of expertise in networking, as well as with researchers in other disciplines (e.g., user experience).¹

2 Measuring QoE Assessment of Real Users

Researchers tend to assess the QoE of users in one of two ways: (1) through controlled laboratory experiments where the researcher varies network conditions and tracks the user's quality of experience by soliciting explicit feedback; (2) through field experiments, where the network performance varies *in situ* and the user's quality of experience is inferred through indirect measures, explicit feedback, or experience sampling.

Laboratory Experiments. There has been a significant amount of laboratory research tracking the performance experience by users interacting with specific applications while they are subjected to various networked conditions. These experiments provide user response and feedback about application performance based on a wide range of network parameter settings. It is also well-understood how to conduct and normalize the quality of experience scores that a cohort of users may report when all users are subjected to the same experiment. Although a randomized control experiment in the lab may in principle provide a more solid conclusion, assuming proper scale and good controls, the cost of running an experiment that yields an appropriate sample can impose barriers to progress.

Field Experiments. Several presenters in this session suggested the value of moving from highly controlled lab experiments to larger real-world experiments. In field trials, the researcher assesses the user's quality of experience either indirectly (e.g., if the user abandons the application) or through explicit feedback (e.g., surveys, experience sampling, feedback buttons). The proliferation of device types, applications, and contexts presents an opportunity to explore user QoE on a much larger scale, albeit with a non-probabilistic sample. Field experiments may include *natural experiments*, whereby individuals experience network conditions as they naturally vary. In some cases, one can argue that the distribution of the experimental conditions approximates random assignment [1]. A slightly more controlled variant is a *quasi-experiment*, where the experimenter has some control over the experimental variables, but the assignment of users to different experimental groups is not random [2]. The concepts of *natural experiment* and *quasi-experiment* have been developed in the social sciences and medical research. The networking community is still determining the best way to apply these measurements.

One opportunity for field experiment at massive scale arises in the context of popular platforms for content delivery. Instrumented client software can observe millions of user interactions, which can (with sufficient capture of resource and contextual variables) provide a basis for relating QoE to observable indicators [3]; these methods, however, typically require incredibly large scale. Experiments that depended on user feedback had to run for a long time to get enough samples without irritating the subjects with frequent query interruptions [4]. Field experiments with mobile devices are often more complex because active network measurements may consume data caps or expend energy, and because network performance (and ultimately QoE) depends - perhaps more than in the wired context - on cross-layer interaction [5].

Assessing QoE in the Field. Indirect measurements may sometimes yield noisy results given the wide range of factors, many outside the reach of the experiment, that determine users' QoE. For example, while user abandonment of an application like video on demand may be indicative of poor QoE, abandonment could also be triggered by other contextual elements such as other activities surrounding the user (e.g., a door ring). Much has been learned about what influences the way users report QoE, such as the fact that the quality of the experience just before a reporting event has a dominant influence or that a sudden loss of quality is more significant than an improvement. Some weak signals must be measured at very large scales to achieve confidence and many indirect measures also require some understanding of human psychology (e.g., using abandonment as a metric requires understanding the factors that may affect a user's level of patience).

3 Correlating QoE to Network Measurements

Mapping QoE to network-layer metrics. Considering the challenges of assessing QoE, a natural followup question involves whether these QoE metrics can be reasonably map to network metrics that could more easily be gathered a automated fashion. To this end, one of the presenters discussed the use of multiple machine learning algorithms for prediction QoE based on networks and systems metrics, and their results in the context of VoIP and VoD on wireless. While most previous work has focused on modeling has targeted an "average user", their proposed ML-based approach could be used to capture an individual user's preferences [6].

Key to this approach is the collection of application-specific and QoE metrics. Another one of the speakers [7] proposed the idea of in-app toolkits, for crowdsource the collection of such measurements, and the instrumentation of WiFi access points. This was motivated in part by the need to measure at different layers, observing correlations and trying to derive causality; as

¹This summary report describes the topics discussed and summarize the key points of the discussion. All materials related to the workshop, including cited white papers, are available at <http://aqualab.cs.northwestern.edu/conference/276#program>

well as the need to engage application developers in design discussions concerning the extraction of various QoE metrics from the applications themselves.

Challenges posed by network-layer encryption. The proliferation of end-to-end encryption can in many cases make it more challenging to extract network-layer metrics, which may in turn make it difficult to predict QoE from these metrics. One of the presenters in this session [8] discussed the complications that the increasing use of encrypted video traffic brings to the ability of network operators to monitor and manage network performance (*i.e.*, the inability to rely on DPI to extract information from a traffic flow beyond the TCP/IP headers). The proposed method for network-wide QoE monitoring works with encrypted traffic, using standard radio network statistics and protocol header information. The discussion touched on ways to improve QoE without breaking end-to-end encryption model (perhaps by using some bits up in the stack), the kind of information needed by operators, and possible collaborations between the main stakeholders: CDNs, users and network providers.

4 Application Design

Another research area concerns how the design of applications themselves can ultimately affect QoE. Several speakers discussed how application design can affect QoE. Assuming that the application designer has tools to assess the current QoE, the app can take adaptive steps to compensate for the assumed cause of the impairment. This space is large and application-specific.

Evaluating application designs with field experiments. In field experiments, application designers can regularly make changes to applications and observe how users' QoE varies. Much application design is focused on improving QoE, since such improvements are essential to user adoption and success in the market [9, 10]. For an application such as video streaming, the options for adapting the delivery are complex, including cache placement, the network used to feed the cache, source selection, and adaptive coding. Content providers may use multiple CDNs to further enrich their delivery options [10]. Yet, it was also noted that highly variable conditions often correspond to a lower reported QoE. To this end, the stability of quality would seem to be an important metric, possibly even more important than the absolute quality itself. More research is clearly needed to understand this phenomenon.

QoE libraries for application developers. A possible aid to building QoE-aware applications would be a set of libraries that could be included in applications to ease the challenge a new application designer faces in measuring QoE; it is an open question as to whether such libraries could even exist, given that different applications may have different notions of QoE. It was noted that there have been some past efforts in this area—notably Web 100, which exposed TCP state variables to the application; and the Android API, which exposes some QoE metrics to applications. It was noted that the Web 100 API in particular has never been heavily used by any applications.

5 Cross-layer Adaptation and Information Sharing

Applications can be designed to adapt in the face of QoE impairments that result from changes in network conditions. For example, it is common for many streaming video applications to change the bitrate of the encoded video mid-stream to adapt to changing conditions. When a client experiences performance degradation, for example, it might request the next video chunks at lower bitrate encodings.

Sharing information across layers. One proposal involved bi-directional information flow that would allow not only the application to adapt to changing network conditions, but also the network layer to adapt to changing application requirements [11]. In this sense, information about conditions can flow in both directions: the network to applications, and vice versa. The discussion included various examples of cross-layer signaling in both directions (*e.g.*, ECN, MTG, SPUD, MPEG-DASH). The discussion noted that because content providers and ISPs each have a different view of the network stack, there is an inherent tussle between different stakeholders concerning information sharing. Another challenge is that for such APIs to exist, they must be common across applications and hence agnostic to the requirements of any particular application. Other challenges included the need for cooperation across multiple ISPs along a path (*e.g.*, to agree on a common signaling protocol), potential security challenges, and the fact that incentives to divulge such information are not necessarily aligned across all stakeholders.

Incentives for cross-layer sharing. To permit this type of signaling in practice, the parties that would be involved must have the incentives to cooperate. Sharing of data among actors will always raise issues, since it might happen under certain circumstances that what is revealed is not in the best interest of one or another actor. Therefore, an incentive structure that favors cooperation must be designed along with the technical features of the proposal. One of the presenters [12] posed the question of the extent to which cooperative management schemes and underlying business models with multiple players can achieve efficient management of network/system resources. There was disagreement among the participants as to whether different stakeholders necessarily share these incentives to share such information. One participant said that content providers had no incentive to share; another said that there was evidence of a willingness to share but the lack of protocols and mechanisms was a major barrier. It was claimed that at least bi-lateral sharing is happening today; one person suggested that game theory supports the claim that there are valid incentives to share information.

6 Challenges in Measuring Network Performance (QoS)

Measurement of network performance (QoS) is a necessary component of understanding and improving QoE. Speakers discussed a variety of approaches to network measurement that might shed light on whether there are impairments that would affect QoE. The approaches included probing to detect variation in latency across links in the core of the network (suggesting queues of traffic due to congestion) and looking for signatures in the variation in round-trip delay as a hint as to where congestion is occurring. If traffic encounters a link that is already congested, its packets may well encounter a standing queue of traffic, which would imply higher delay but small variation in delay, while a flow that induces its own congestion (for example on an access link) will show a much higher variation in latency as the congestion builds up.

Challenges introduced by middleboxes. Measuring the network is becoming more difficult as the network itself becomes more complex. Middleboxes may perform complex transformations on data being sent, so that what comes out at the application layer is not what goes in. Attempts to measure the network using a simple data transfer (e.g., an end-to-end TCP transfer) may not give a reliable indication of what would happen to a data flow in a specific application, while on the other hand the transformation of the application data by a middlebox may mislead application level measures of network performance [13].

Challenges in measuring congestion and its effects on QoE. Different network paths may encounter different path conditions: for example a path through a public peering point may see different congestion than a path through a private interconnection link. Networks may do routing based on inspection of the data being sent, again confounding simplistic measurements. And again this landscape may change rapidly with the movement toward encryption. It may be appropriate for the research community to consider whether there are other sorts of service indicators (e.g., new values of the DSCP field) that could trigger new sorts of service adaptation by the network [14].

Performance measurement at IXPs and private interconnects. Measurements at IXPs can help determine how capacity should be provisioned at interconnection points [15]; for example, QoE measurements can help determine whether more virtual private interconnection ports should be provisioned to exchange traffic. Determining how the performance at both public IXPs and private interconnects should be measured is an open research problem. Specifically, because continuous active monitoring is so expensive, some IXP operators are exploring the use of active measurements that are triggered by particular events, such as a network performance anomaly. Many large CDN providers perform a combination of both active and passive measurements at these interconnects, but the design of these measurement strategies remains an open problem.

Another open problem involves how data that is collected at interconnection points can be constructively shared with data collected at other vantage points. For example, passive measurements collected at an IXP or private interconnect could be analyzed in tandem with active measurements collected at the network edge (as in the previously discussed active probing work [14]). Yet, passive data at private interconnects is rarely shared—and, when it is shared, the data is often released in coarse aggregates. Developing methods to combine passive data collection at interconnect points with QoE measurements from clients or network performance measurements from the edge of the network is another research challenge.

7 Economics of QoE

As the speakers in this session emphasized, economic considerations can ultimately influence QoE.

Interconnection agreements and QoE. Certain aspects of the network topology, such as which networks connect with one another and how much capacity is provisioned between networks, is often determined by the interconnection agreements between ISPs and content providers [16, 17]. Certain types of interconnection arrangements increasingly see high utilization—which, in some cases, causes congestion that results in degraded QoE. Due to the shortcomings of current network measurement techniques (as previously discussed in this report), it can be difficult to determine how the utilization levels at interconnection points ultimately affect the QoE of applications for end users. Certain application providers may also drive utilization to higher levels, which may result in degraded QoE in some cases. If QoE depends on properly provisioned links, an important question remains: Who should pay for the provisioning of these links? Ultimately, the economics of these business relationships and how these types of peering disputes are (or are not) resolved may ultimately have significant effects on QoE.

Relationship between QoE and zero-rated services. Users have a higher expectation for services that they pay for, and QoE is reported relative to expectation. The relationship between economics and QoE also entails more complex considerations. Users that are about to exhaust their monthly quota may show very different behavior than those not concerned about hitting their limit. Users downloading content that does not affect their monthly quota (*i.e.*, so-called “zero rated” content) may have a different expectation for QoE than if the application was consuming their allocated data quota [18].

8 Localization and Coordination

To detect that there is an impairment to QoE does not give much guidance as to how to remedy it. A necessary step is to localize the problem. Coordination among several sorts of measurement may be necessary to achieve that localization and identification of root cause.

Coordination between users and ISPs. As discussed above, different users will have different views of application, resource, and context factors that influence QoE. Users may wish to exchange information with their ISPs about the QoE that they are experiencing. Additionally, coordinated active probing may help locate the place in the network where QoE is impaired, but there is a risk that certain networks may distort active probing measurements (*e.g.*, by prioritizing network probe traffic). Comcast’s current efforts toward building predictive models of QoE are based on various metrics of health for different portions of the network [19]. Target metrics for these models include self-reported issues (*e.g.*, logs of service calls or opinion surveys) or directly observed behavior. One of the participants pointed out some of the challenges related to these efforts, such as putting together and making sense of the dataset, from aligning the times of adverse network events and customer calls, to understanding relevant contextual (*e.g.*, time of day) and confounding factors, as well as the impact of history (*e.g.*, a user may not call the first time they experience a service degradation, but there are more likely to do it the third time around)

Coordination between ISPs at IXPs. Recent advances in programmable hardware switches (*e.g.*, P4’s in-band network telemetry) may soon make it possible to design mechanisms for measuring application performance in-band, as packets for particular flows traverse the network [20]. Specifically, recent technology such as in-band network telemetry will ultimately make it possible for network routers and switches to annotate network traffic with fine-grained information about per-hop latency and packet loss through the network. If deployed at Internet exchange points (IXPs), this new capability may ultimately make it possible to gather fine-grained, per-flow, per-hop information about individual flows in the network. Whereas capturing fine-grained information about latency, buffer occupancy, and packet loss is difficult to implement at high-speed interconnects today, emerging technology suggests that these capabilities may soon be a reality, thus creating possibilities for considerably more information about the network performance (QoS) at specific points along an end-to-end path, including at interconnection points. In addition to *monitoring*, these programmable switches may also provide the possibility for *control*, such as implementing reactive rate-shaping of high-rate flows, to preserve the QoE for all flows that are traversing the interconnect.

Coordination between users. One speaker described a scheme that allow different users to collaborate in order to see whether they are having a similar experience or one that is unique to a single user. This approach, which does not require the cooperation of the network providers, can in practice provide a effective means for at least partial localization [21].

Finally, one of the presenters proposed an approach to coordinate the flow of QoE information from end users, apps, or devices to access providers, content and service providers, hardware manufacturers, and software developers [22].

9 Measurement infrastructure

The workshop also included a discussion, led by K.C. Claffy, of the type of measurement infrastructure that could improve the measurement of QoE on the Internet. In addition to researchers who can benefit from better visibility into network performance, both network operators and application developers can benefit in detecting and correcting problems with their systems. If the infrastructure includes instrumentation embedded in network and in client-server system and application code that is ultimately used by these stakeholders, such a system could be a resource that is shared by a wider community, including researchers, developers, and operators.

Workshop participants discussed the following questions:

- What results might come from various forms of collaboration?
- How can different sorts of data be combined and correlated to help identify the locations where QoE impairments are occurring?
- Do we need new sorts of measurements?
- Are there specific research efforts that we should identify?

Participants also completed a long-form survey where they addressed some of these questions. The rest of this section summarizes some of the findings from that survey.

QoE measurement infrastructure for general applications. Several respondents suggested that the infrastructure include the capability to measure client-and server side application performance—the key performance indicators (KPIs) that signal QoE, as well as network performance. As a complement to measurement, we need to advance the state of understanding about models that relate user-level QoE to key performance indicators (KPIs) that can determine what lower-level measurements are needed.

There was agreement on the need for several types of measurement, both active and passive, background (always running) and on-demand (fault diagnosis). Participants also recognized the need for measurement infrastructure that could better capture higher-level performance metrics, such as those that more directly focus on user experience. Several ideas on these topics

were discussed including the use of virtualized infrastructure that could provide “QoE measurements as a service”. Given the application-specific nature of QoE, as well as the many end-to-end network components that ultimately affect QoE (*e.g.*, the device, the home network, the ISP, interconnection, server provisioning, the application), it remains unclear whether it is possible to build a common infrastructure to reflect the general state of QoE on the Internet.

Correlation of heterogeneous signals. The traditional conception of a measurement infrastructure is a set of devices that can perform measurements as required. However, several of the respondents noted a second requirement, which is a system that can gather, aggregate and correlate different sorts of measurements to provide an overall assessment of a specific property such as QoE, or perform fault diagnosis using this aggregation of information. We (as a community) have much less experience with this sort of system. Even with open platforms (such as PlanetLab), the construction of an overall system has been the purview of individual researchers. Would such an “aggregation infrastructure” (as opposed to a “measurement infrastructure”) be the primary goal of a significant research effort? Would it be an end-goal or yet another platform for use by another tier of researchers? To what extent would these two platforms be integrated?

Several of the respondents noted that most datasets remain disjoint, and joint analysis and correlation of these disjoint datasets remains difficult, given the infrastructure that is available to the community. Many researchers believed that the participation of ISPs in standing up this infrastructure was critical to having a successful infrastructure, as many of the vantage points for the data would likely be from within those ISPs. Another suggestion was to look at the current error reporting systems developed by OS and application developers; these systems are providing real value and the network community is far behind in this respect. (Note—error reporting systems are one way to get actual user feedback and signals of impaired QoE from the user perspective.)

Who will lead the effort to create this infrastructure, and who will own it? An issue that was both mentioned explicitly and implicitly signaled in other comments was the need to think hard about the incentives that would make this infrastructure come into existence. To encourage the integration of measurement into applications, it would be good to develop a measurement library, preferably open source. The design of such a library would be an important, but significant effort. In particular, to encourage the instrumentation and measurement of home networks (considered to be a source of many impairments), the users must have a suitable incentive, most obviously that the data benefits the users themselves—it helps them to understand and correct flaws in their networks. In the design of a large system such as this, one of the problems to solve is the “first-mover disadvantage”. The real benefit emerges only after a number of actors have undertaken to participate, so the first movers incur cost but at first gain little benefit.

NSF-funded development of a measurement platform for the community would suggest that the infrastructure be “open”, to the extent possible (*i.e.*, various researchers should be able to use it for their experiments). However, there are several risks associated with an open platform. A platform that does passive measurement (observes ongoing traffic) raises serious privacy risks. A platform that performs active measurement, on the other hand, can overload the network, consume access quotas (for devices in the home), and risk abuse complaints. Beyond the measurement and data collection infrastructure itself, the justification for NSF support of the infrastructure that supports correlation and aggregation analysis across large, distributed, and heterogeneous datasets is even more complex, since development of such an analysis framework may require new research.

10 Closing Remarks and Next Steps

The Internet, for most people, is a commodity. We have to get to the point where it just works; providing adequate QoE is a critical, indispensable part of that story. Attendees concluded that a second workshop was warranted, but that further work was required to define the right focus and objective.

- With respect to the controlled measurement of QoE, there is a substantial body of recent work, much of it done in Europe.² Having yet another workshop on this topic without first understanding the current state of affairs is unlikely to be valuable.
- With respect to application design to improve QoE, there were few suggestions for application-independent approaches to this problem. It is not clear how to organize a workshop around this set of issues, but there may be general lessons that could be extracted.
- With respect to cross-layer optimization, there is significant work going on in the mobile/wireless context, and a better exploration of that work would be a useful foundation for a follow-on workshop. One issue that might warrant explicit attention is incentives, and how to motivate cross-layer sharing.
- With respect to network measurement and infrastructure for QoE measurement, aggregation, and analysis, it was decided to exploit CAIDA’s Active Internet Measurement Systems workshop to explore that issue further. That workshop will produce a final report that will further inform this work.
- With respect to localization, the workshop uncovered many open challenges, ranging from defining new protocols for measurement to the potential for coordination between users and ISPs, applications and ISPs, or between the ISPs

²A good summary of the current state of knowledge can be found in *Quality of Experience: Advanced Concepts, Applications and Methods*, Springer, 2015.

themselves. Emerging hardware capabilities may also facilitate measurement at interconnection points that could ultimately be revealing. These topics appear ripe for further study.

References

- [1] Fabián E. Bustamante. Natural Experiments for Quality of Experience in Networked Systems. White-paper submission, Northwestern University, 2015. (Cited on page 2.)
- [2] Ramesh K. Sitaraman. Internet Performance: Who Cares and by How Much? White-paper submission, U. of Massachusetts, Amherst/Akamai, 2015. (Cited on page 2.)
- [3] Aditya Ganjam, Vyas Sekar, and Hui Zhang. In-Situ Quality of Experience Monitoring: The Case for Prioritizing Coverage Over Fidelity. White-paper submission, CMU/Conviva, 2016. (Cited on page 2.)
- [4] Renata Teixeira. Measuring Networks for Enhancing User Experience (MUSE). White-paper submission, Inria Paris-Rocquencourt, 2015. (Cited on page 2.)
- [5] Z. Morley Mao. Addressing QoE Measurement Challenges for the Mobile Internet. White-paper submission, U. of Michigan, 2015. (Cited on page 2.)
- [6] Maria Papadopouli. On User-Centric QoE Prediction for VoIP and Video Based on Machine-Learning. White-paper submission, U. of Crete/ICS-FORTH, 2015. (Cited on page 2.)
- [7] Suman Banerjee. QoE: A Wireless Perspective and Beyond. White-paper submission, U. of Wisconsin, Madison, 2015. (Cited on page 2.)
- [8] M. Zubair Shafiq. Tracking Mobile Video QoE in the Encrypted Internet. White-paper submission, U. of Iowa, 2015. (Cited on page 3.)
- [9] Yong Liu. Assessing and Improving User QoE of Video Telephony. White-paper submission, New York University, 2015. (Cited on page 3.)
- [10] Zhi-Li Zhang. Complex Interactions in the Content Distribution Ecosystem and Quality-of-Experience. White-paper submission, University of Minnesota, 2015. (Cited on page 3.)
- [11] Srikanth Sundaresan and M. Zubair Shafiq. Bidirectional Cross-Layer QoE Optimization. White-paper submission, ICSI and U. of Iowa, 2015. (Cited on page 3.)
- [12] Lea Skorin-Kapov. Tracking and Managing QoE in the Internet: Towards Cooperative QoE Management Schemes. White-paper submission, U. of Zagreb, Croatia, 2015. (Cited on page 3.)
- [13] David Choffnes, Phillipa Gill, and Alan Mislove. A Practical System for Revealing ISP Traffic Differentiation. White-paper submission, Northeastern U. and Stony Brook U., 2015. (Cited on page 4.)
- [14] Amogh Dhamdhere, Matthew Luckie, David Clark, and kc claffy. Measuring Interdomain Congestion and its Impact on QoE. White-paper submission, CAIDA/UC San Diego and MIT, 2015. (Cited on page 4.)
- [15] Walter Willinger, Anja Feldmann, Philipp Richter, Georgios Smaragdakis, and Fabián E. Bustamante. Express or Local Lanes - On Assessing QoE over Private vs. Public Peering Links. White-paper submission, NIKSUN and TU Berlin and Northwestern U., 2015. (Cited on page 4.)
- [16] Constantine Dovrolis. The Issue of Poor QoE Viewed as a Problem of Bad Peering Economics. White-paper submission, Georgia Institute of Technology. (Cited on page 4.)
- [17] William Lehr. QoE Measurement: Economic and Policy Considerations. White-paper submission, MIT, 2015. (Cited on page 4.)
- [18] Matthew Andrews. Quality of Experience in Wireless Pricing: Some Research Directions. White-paper submission, Alcatel-Lucent Bell Labs, 2015. (Cited on page 4.)
- [19] Yana Kane-Esrig. Need for Methodologies to Optimize Time-Alignment of Predictor Variables in Building Models for Quality of Experience as a Function of Network Behavior. White-paper submission, Comcast, 2015. (Cited on page 5.)

- [20] Nick Feamster. Monitoring and Measuring QoE in the Wide Area. White-paper submission, Princeton University, 2015. (Cited on page 5.)
- [21] Henning Schulzrinne, Kyung-Hwa Kim, Hyunwoo Nam, Vishal Singh, and Daniel Song. QoE: A Micro and Macro Perspective. White-paper submission, Columbia University, 2015. (Cited on page 5.)
- [22] Steve Bauer. Net.info: A Provider to Subscriber Secure Communicatio Channel. White-paper submission, MIT. (Cited on page 5.)