

Infrastructure to support QoE measurement

kc claffy
CAIDA / UC San Diego
kc@caida.org

Quality of Experience (QoE) is a measure of whether a use of the Internet meets the needs and (“reasonable”) expectations of the user and provider in the context of any specific application. But QoE is a subjective measure of quality. While *quality of service (QoS)* is used to describe the technical parameters of any service (peak speed, latency, jitter, packet loss), assessment of QoE is application-specific. Impairments to QoE derive in part from underlying QoS parameters (e.g., jitter impairs voice QoE but not email), and can arise anywhere in the network.

CAIDA and MIT have demonstrated a research method that can detect such congestion at points of interconnection [1], and found that business contention (usually well-reported in the press) correlates with evidence of substantial persistent congestion at points of interconnection – some episodes have persisted for days or months. Our current research method focuses on measuring persistent rather than instantaneous congestion, but with instrumentation and refinement of such methods, the research community could develop tools that they, the consumer or the FCC could use to gather evidence on whether jitter is arising at points of interconnection or elsewhere in the network.

However, a comprehensive assessment of congestion and its impact on QoE will require new ways of thinking not only about measurement research, but about measurement infrastructure and its use by the research community. As an example, we have encountered daunting challenges when trying to effectively apply our research method at Internet scale: identifying congestion on links with AQM and WFQ policies; accurately finding and identifying all interdomain links involving the AS hosting a vantage point; proving the response from the far router returns over the targeted interdomain link; determining the direction of congestion; robustness to ICMP queuing behavior; adapting to path dynamics; and scaling processing to thousands of interdomain links. These problems require large scale measurement infrastructure specifically architected to accommodate macroscopic assessments, and ultimately we also need feedback from ISPs to establish the validity of our methods for inferring congestion.

For years, the FCC has been inviting researchers to deploy measurement capabilities on the Measure Broadband America program, but the architecture of the measurement platform acutely constrains how researchers can use it. It is not research infrastructure. Google’s Measurement Lab infrastructure was geared to support a certain type of measurement (clients sending active probes to measurement servers), and has policy constraints that limit the research commu-

nity from using it to undertake a more comprehensive assessment of congestion on the Internet. NSF has funded several Internet measurement research infrastructures (including CAIDA’s Archipelago, which is supporting our congestion research, among other projects), but has not thus far undertaken a concerted effort to determine how these infrastructures could work together, or separately, to support a grand challenge of measuring and analyzing the state of QoE in the Internet.

For example, there are infrastructure requirements on both sides: server side infrastructure to which users can run tests (currently MLab is the only such infrastructure that is open); and client side software and/or hardware infrastructure from which to run them. There are also various end-host client system infrastructures, but they are not integrated in a way that facilitates their use for coordinated assessments.

There are also systems challenges: building large-scale systems that can collect and process many diverse types of performance and perception data and make sense of it. The effort of building such systems usually does not look attractive from the research perspective but is critical to enabling QoE research, not to mention sustaining an operational capability to evaluate current performance impairments.

There is an even a more basic question: do we have clarity on what type of infrastructure we need to measure QoE? One thing I hope to get out of the workshop is an “infrastructure/systems wish list” for this line of research? What do we wish we had in this space that would make QoE research easier (and less tedious, more fun) to do, or enable research that is not possible now? Developing a wish list like that from workshop attendees would be quite useful for those in the research community trying to evolve and adapt their own infrastructure, and would inform NSF infrastructure funding trajectories.

Based on our years of experience designing, developing, operating, evolving, and using measurement infrastructure for the Internet research community, my hope is that as we try to cohere a set of research activities that could generate a new way to tackle the problem of measuring QoE on the global Internet, let us please consider the infrastructure challenges in parallel with trying to understand the research landscape. Not after.

1. REFERENCES

- [1] Matthew Luckie and Amogh Dhamhere and David D. Clark and Bradley Huffaker and kc Claffy. Challenges in Inferring Internet Interdomain Congestion. In *Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)*, 2014.