

# Internet Performance: Who Cares and by How Much?

Ramesh K. Sitaraman\*<sup>†</sup>

University of Massachusetts, Amherst & Akamai Technologies

Email: ramesh@cs.umass.edu

## 1 A Personal Journey from QoS to QoE and Beyond

Internet services are at the core of human activity today. Billions of users interact with networks of various kinds for commerce, entertainment, news, and social networking. But, in the late 1990's, things were different. E-commerce was in its infancy. Streaming video was an euphemism for audio, as Internet radio dominated the scene. Social networks didn't exist. And, at Akamai, we were creating one of the world's first content delivery networks (CDNs) [3].

All networks, including CDNs, are designed with a notion of performance in mind, a notion commonly referred to as *quality-of-service (QoS)*. For a content delivery service, good QoS means web pages download faster and videos start up quickly and play without rebuffers (i.e., freezes). As computer scientists, we understand how to architect networks to provide good QoS. After all much of computer science is a study of performance in one form or another. But, we do not have a good scientific grasp of the relationships between QoS, the more subjective notion of *quality of experience (QoE)*, and the *causal* impact of QoS on QoE, user behavior, and economic outcomes. In this paper, I argue that the time is ripe for unraveling these relationships, a research agenda that could lead to vastly better network architectures, deeper knowledge of user behavior, better monetization, and more informed Internet policies and governance.

### 1.1 The Economic Life of Networks

My first brush with the complex relationships between QoS, QoE, user behavior, and economic outcomes occurred more than fifteen years ago when we journeyed into the world of content providers to encourage them to use CDNs that we had just built, a relatively new and unfamiliar technology at that time. Content providers, rightfully focused on their business, would often ask questions for which I could find no *rigorous* answers: Would faster web page downloads cause users to buy more from my site?

---

\***Disclaimer:** Opinions expressed in this paper are solely that of the author and not necessarily that of Akamai or UMass or any other organization.

<sup>†</sup>Supported in part by NSF grant CNS-1413998.

Would fewer rebuffers result in viewers watching more videos? When would users abandon videos that are slow to load? Would users return to my web site less often if they experience sporadic failures? These are not questions about a CDN's QoS per se, but about how QoS impacts the user's QoE and behavior, and how that in turn impacts the content provider's business metrics. Answering such questions with rigor has been a preoccupation of my research over the years.

**The virtuous cycle.** The causal relationships between QoS, QoE, user behavior and monetization form what we call a virtuous cycle (cf. Figure 1) [4]. The virtuous cycle is central to the economic life of a real-world network service and captures the interactions between three entities: an online business, their users, and the network service provider. A business invests in a network service to provide better QoS for its users. Better QoS may lead to a better QoE, impacting the manner in which users interact with the business. In turn, favorable changes in user behavior may lead to greater profits for the business. The increased profits justify current and future investments in the network service, closing the cycle.

As an concrete example of a virtuous cycle, content providers may invest in a CDN to deliver videos to their users with better QoS in the form of smaller video startup delays and fewer rebuffers. The expectation of the content providers is that better QoS would lead to better QoE, that will in turn alter the viewing habits of their user population in ways that are favorable to their business, e.g., the fraction of viewers who abandon their videos may decrease, the amount of time each video is watched may increase, and the likelihood that viewers return to their web site over time may become greater. Since content providers typically rely on advertisement or subscription revenues, more opportunities for ad insertion and a larger viewership base mean more revenues. Thus, all three goals of reducing abandonment, increasing engagement, and improving repeat viewership lead to increased monetization of their business, that in turn justifies their original investment in the CDN.

**The Missing Link: Impact of QoS on QoE and User Behavior.** The key to designing a better real-world network is to understand its virtuous cycle.

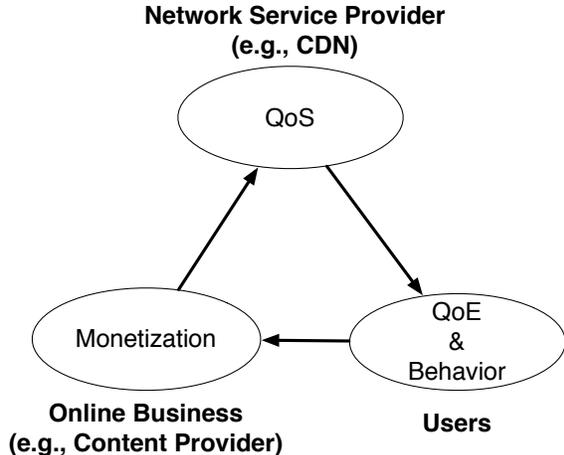


Figure 1: The virtuous cycle links QoS, QoE, user behavior, and monetization.

For instance, one cannot design a good video CDN without understanding how the two opposing QoS requirements of small video startup delay and reduced rebuffering impact user behavior and monetization. QoS, QoE, user behavior, and monetization are studied in isolation today, often within entirely different disciplines, without an understanding of how all the pieces fit together. The link between user behavior and monetization is often well understood. However, the missing link is a *scientific understanding of how QoS influences QoE and user behavior*.

**Quantification and Causality.** The fact that better QoS would have a positive impact on users would seem almost tautological. So, why study this at all? The key reason is that any study of the impact must have two characteristics to be useful: it must be *quantifiable* and it must establish *causality*. Taking video delivery as an example, suppose we hypothesize that video rebuffering causes users to play fewer minutes of the video. There is certainly value in establishing the hypothesis as is. But, it is more valuable to show *quantitatively* that an  $x\%$  increase in video rebuffering causes a  $y\%$  decrease in the minutes of video watched. A quantitative statement of this sort enables us to estimate the potential user impact of reducing rebuffering and determine if it is worth the effort. Further, the impact of QoS on the user must be established in a *causal* manner. The virtuous cycle is one of cause and effect. A content provider investing in a network service needs to know that his/her investment in improving QoS was the *cause* for the observed changes in viewer behavior, and not just something that would have happened anyway for other reasons.

## 1.2 Recent Progress

For much of the past fifteen years, reliable large-scale measurements of QoS, QoE, and user behavior was

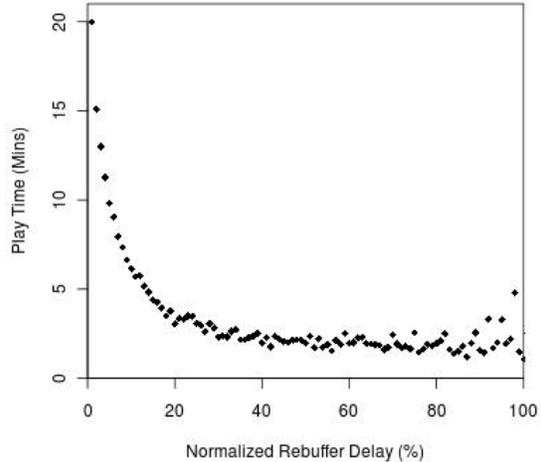


Figure 2: Users who suffer a higher rebuffer delay play fewer minutes of the video [1].

hard to come by. But, now it is possible to measure both user-perceived QoS and user behavioral actions at a large scale from actual users in the field. For instance, in the case of online videos, it is now possible to integrate an analytics plugin that runs inside the user’s media player that measures and reports both user actions (play, pause, rewind, browser close, etc) and QoS metrics (startup delay, rebuffering events, bandwidth, etc). Our recent research [1, 2] rely on such large scale measurements from hundreds of millions of video sessions from actual users around the world. Likewise, such research is also possible for web delivery, as modern browsers have standard APIs that can extract QoS and user actions that can be collected and analyzed at a large scale.

Large amounts of data is not enough. We need powerful tools for designing the right experiments that could yield causal answers, taking into account numerous confounding factors that could obfuscate the outcomes. Here we have adapted classical techniques from the social and medical sciences for networking research. In particular, in [1], we introduce quasi-experimental designs (QEDs) to establish a causal and quantitative link between video QoS and user behavior. For instance, we show that video rebuffering causes users to watch fewer minutes of the video (cf. Figure 2) and also that if a video fails to load the user less likely to return to the same site.

## 2 Science-Informed Policy: User Behavior and Net Neutrality

Quite apart from the economics of networks, understanding the behavior of humans interacting with online services is important in its own right from per-

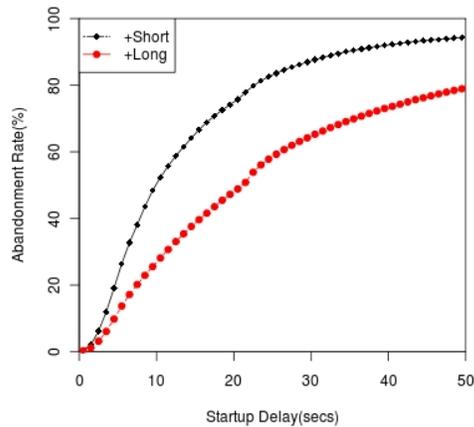


Figure 3: Users abandon videos at a higher rate for short-form videos than for long-form videos [1].

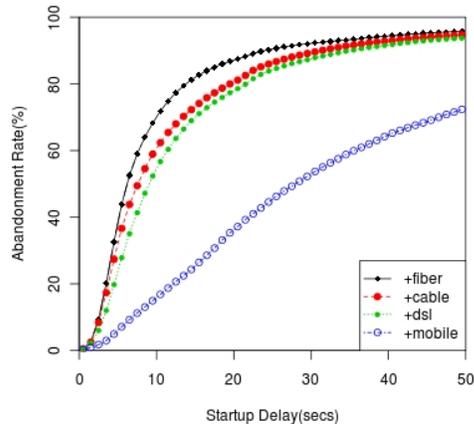


Figure 4: Users who have better Internet connectivity (say, fiber) have less patience for the video to start up and abandon at a greater rate than those with slower Internet connectivity (say, mobile) [1].

spectives that range from plain scientific curiosity to better informed Internet governance and policies. Hence, exploring the principles of online human behavior is a thrust of my research, an example below.

**How humans wait in the online world?** Much is known about how humans wait for services in the physical world. The “psychology of waiting” as it is often called is widely applied wherever humans wait in the physical world: from Disney World to your local grocery store. But little scientific is known about human patience and waiting behavior in the online world. Using quasi-experiments, we recently studied how long users are willing to wait for their video to start playing. Computing the *abandonment rate* which is the rate at which users abandoned their video while waiting for it to start up, we showed that users

wait longer for long-form videos such as movies than for short-form videos such as news clips [1]. Further, users are more patient and abandon less often when watching a pre-roll ad, than when they wait watching nothing but the “spinning globe” [2].

Using quasi-experiments, we also studied the influence of Internet speed on user patience [1]. We showed that users with faster Internet connectivity have greater expectations of speed and thus had lesser patience to wait for a slow-loading video to start (cf. Figure 4). Specifically, users with faster connectivity (e.g., fiber) abandoned a slow-loading video at a greater rate than users with slower Internet connectivity (e.g., mobile). Our results provide *quantitative* evidence that the patience level of users is influenced by the expectation of Internet speed, a fact of key significance to the recent net neutrality debate. Our results also suggest that a “fast lane” in the Internet can decrease the user’s tolerance to the relative slowness of a “slow lane”, making it harder for online businesses on the slow lane to compete with those who can afford the fast lane. For this reason, these results [1] were widely cited in defense of net neutrality by content providers such as Vimeo and Netflix.

### 3 Concluding Remarks

Understanding the interplay of QoS, QoE, user behavior and monetization is an important area for research. The fruits of such research could lead to better network architectures, deeper knowledge of user behavior, better monetization, and better informed Internet policies and governance.

### References

- [1] S Shunmuga Krishnan and Ramesh K Sitaraman. Video stream quality impacts viewer behavior: inferring causality using quasi-experimental designs. In *Proceedings of the ACM Internet measurement conference*, pages 211–224. ACM, 2012.
- [2] S Shunmuga Krishnan and Ramesh K Sitaraman. Understanding the effectiveness of video ads: a measurement study. In *Proceedings of the ACM Internet measurement conference*, 2013.
- [3] E. Nygren, R.K. Sitaraman, and J. Sun. The Akamai Network: A platform for high-performance Internet applications. *ACM SIGOPS Operating Systems Review*, 44(3):2–19, 2010.
- [4] Ramesh K Sitaraman. Network performance: Does it really matter to users and by how much? In *Fifth International Conference on Communication Systems and Networks (COMSNETS)*, pages 1–10. IEEE, 2013.