

# Tracking Quality-of-Experience in the Internet

Constantine Dovrolis (Georgia Tech)

QoE is hard to define in any formal manner. Part of the issue is that it relates to many different things: the user herself (experience, expectations, psychological state), the application’s requirements and implementation, the way the application interfaces with the network through operating system and middleware mechanisms, the bandwidth and buffer provisioning in the involved networks, the interdomain interconnections and the policy and economic issues that are associated with those, etc. Over the last years we have worked on different problems that relate to different aspects of QoE, focusing mostly on video streaming applications (mostly Dynamic Adaptive Streaming over HTTP), on the impact of WiFi access in QoE, on network probing methods to detect and localize performance problems, and on the economic/policy issues associated with interdomain interconnections (peering and transit).

This white paper first summarizes the research that our group has been conducting over the last 5 years or so in the broader area of Quality-of-Experience (QoE) monitoring. Second, it identifies a couple of open research problems, cross-disciplinary in nature, that we would like to pursue in the next few years in collaboration with others.

## 1 Prior relevant work

**The Internet seems slow today: is it my WiFi?** Particular WLAN pathologies experienced in realistic scenarios are hard to detect, due to the complex nature of the wireless medium. Prior work has employed sophisticated equipment, driver modifications, or even application-layer techniques, towards diagnosing such pathologies. In a series of two publications, we developed probing methods that run at the application-layer, without any access to the WiFi access point, and allow the user to detect if the performance of her WiFi network is problematic. These issues may be caused by low-SNR, congestion, hidden terminals, interference from devices that do not use the IEEE 802.11 protocol, etc. The proposed methods also provide a diagnosis regarding the root cause of the observed problems.

**What is wrong with my Netflix today?** In a series of publications we focused on rate-adaptive video streaming applications that run over HTTP (such as the NetFlix player, among many others). Our first study in 2011 showed that the rate-adaptation algorithms in these video players are inherently unstable because of the way they infer available bandwidth in the network [3]. This instability can cause poor QoE because the video keeps oscillating between different quality profiles instead of converging to the highest quality profile that is sustainable given the underlying network conditions. In subsequent publications we showed what causes this instability in a general setting [1], how to resolve the issue with modifications in the client’s rate adaptation algorithm [4], how to resolve it with modifications in the video server [2], and finally, we showed that similar instabilities can take place under certain conditions when different video fragments are served by different caches/CDN servers [10].

**Is my ISP penalizing some applications?** A major concern in the context of network neutrality is whether ISPs prioritize some Internet traffic, deteriorating the performance of other applications (e.g., serving Over-the-Top VoIP traffic through lower priority queues or rate-limiting BitTorrent traffic). We developed an active probing method, called Differential Probing or *DiffProbe*, to detect whether an ISP is deploying forwarding mechanisms such as priority scheduling, variations of WFQ, or WRED to discriminate against some of its customer flows [8].

**How fast is my Internet connection, really?** There are dozens of Internet “speedometers” available today, such as Ookla’s “speedtest”. Practically all of these tools perform a large TCP transfer from the client to a nearby server, and another transfer in the opposite direction, reporting the average throughput in each direction. This often creates confusion because the average TCP throughput of a single large transfer can be significantly lower than than the capacity of the bottleneck link in an Internet path. Additionally, many ISPs today deploy traffic shaping technologies, such as Comcast’s PowerBoost, that allow the user to receive a higher “peak throughput” for a certain traffic volume, before settling to the nominal capacity of

that link. We have developed an inference method and a measurement tool, referred to as *ShaperProbe*, that allows a user to measure the actual capacity of her Internet connection in both directions, and to detect the presence and parameters of a traffic shaper in the path [9].

**Who should I blame when the Internet is slow?** Most Internet traffic goes through more than one Autonomous System (AS). When a user experiences bad network performance, which network is responsible? In which AS does the congested link belong to? We have relied on network tomography approaches to perform such *localization* of congested Internet links [7, 13]. The objective of early network tomography approaches was to produce a point estimate for the performance of each network link (Analog tomography). When it became clear that the previous approach is error-prone in practice, research shifted to Boolean tomography where each link is estimated as either "good" or "bad". The Boolean approach is more practical but its resolution is too coarse. Our main contribution in this area has been a new tomography framework that combines the best of both worlds: we still distinguish between good and bad links (for practicality reasons) but we also infer a range estimate for the performance of each bad link.

**At the end of the day, poor QoE is often due to bad economics** Especially in the last few years, there is a major focus on the presence of congestion in certain interdomain interconnections and on the reasons that that the involved parties cannot agree on who should pay for the required capacity upgrades. We believe that this issue is mostly of economic, rather than technological, nature and it is caused by the mechanisms through which ASes negotiate and decide their interconnection (transit and peering) agreements. In a series of publications, we have focused on the economics of such interconnections and on the evolution of this peering ecosystem [5, 6, 11, 12]. This research has convinced us that the currently deployed interconnection framework, which is based either on settlement-free peering relations or transit relations, is very limiting in terms of the spectrum of interconnection arrangements that should be feasible, creating situations when none of the involved ASes has the incentive to upgrade the capacity of an interconnection.

## 2 Looking forward

I am committed to continue working in the broader area of Internet QoE monitoring. The specific research directions that I would like to explore in the near future, especially in collaboration with people that have the related expertise, are:

1. Data mining and machine learning methods that can be deployed by either ISPs or Content Providers, allowing them to monitor the QoE of their users based on the massive data they are already collecting about individual flows or sessions. For instance, how can an ISP mine netflow records from its network to automatically detect that a cluster of customers in a specific region/subnetwork experience bad video streaming performance? We have recently started a research project in this area in collaboration with Cisco Systems.
2. We are currently starting a new NSF-funded project, in collaboration with economists, that aims to propose a new techno-economic framework for interdomain interconnections that can provide a broader and more economically efficient set of interdomain relations than just transit and settlement-free peering. The proposed interconnection framework should also provide the right incentives so that all relevant Internet firms continue to invest sufficient resources to Internet infrastructure.

## References

- [1] S. Akhshabi, L. Anantkrishnan, A. C. Begen, and C. Dovrolis. What happens when HTTP adaptive streaming players compete for bandwidth? In *Proceedings of the 22nd international workshop on Network and Operating System Support for Digital Audio and Video*, pages 9–14. ACM, 2012.
- [2] S. Akhshabi, L. Anantkrishnan, C. Dovrolis, and A. C. Begen. Server-based traffic shaping for stabilizing oscillating adaptive streaming players. In *Proceeding of the 23rd ACM Workshop on Network and Operating Systems Support for Digital Audio and Video*, pages 19–24. ACM, 2013.
- [3] S. Akhshabi, A. C. Begen, and C. Dovrolis. An experimental evaluation of rate-adaptation algorithms in adaptive streaming over HTTP. In *Proceedings of the second annual ACM conference on Multimedia systems*, pages 157–168. ACM, 2011.
- [4] S. Akhshabi, S. Narayanaswamy, A. C. Begen, and C. Dovrolis. An experimental evaluation of rate-adaptive video players over HTTP. *Signal Processing: Image Communication*, 27(4):271–287, 2012.
- [5] A. Dhamdhere and C. Dovrolis. The Internet is flat: modeling the transition from a transit hierarchy to a peering mesh. In *Proceedings of the 6th International Conference*, page 21. ACM, 2010.
- [6] A. Dhamdhere and C. Dovrolis. Twelve years in the evolution of the Internet ecosystem. *IEEE/ACM Transactions on Networking (ToN)*, 19(5):1420–1433, 2011.
- [7] A. Dhamdhere, R. Teixeira, C. Dovrolis, and C. Diot. NetDiagnoser: Troubleshooting network unreachabilities using end-to-end probes and routing data. In *Proceedings of the 2007 ACM CoNEXT conference*, page 18. ACM, 2007.
- [8] P. Kanuparth and C. Dovrolis. Diffprobe: detecting ISP service discrimination. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9. IEEE, 2010.
- [9] P. Kanuparth and C. Dovrolis. ShaperProbe: end-to-end detection of ISP traffic shaping using active methods. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, pages 473–482. ACM, 2011.
- [10] D. H. Lee, C. Dovrolis, and A. C. Begen. Caching in HTTP Adaptive Streaming: Friend or Foe? In *Proceedings of Network and Operating System Support on Digital Audio and Video Workshop*, page 31. ACM, 2014.
- [11] A. Lodhi, A. Dhamdhere, and C. Dovrolis. Open peering by Internet transit providers: Peer preference or peer pressure? In *INFOCOM, 2014 Proceedings IEEE*, pages 2562–2570. IEEE, 2014.
- [12] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, et al. Using peeringDB to understand the peering ecosystem. *ACM SIGCOMM Computer Communication Review*, 44(2):20–27, 2014.
- [13] S. Zarifzadeh, M. Gowdagere, and C. Dovrolis. Range tomography: combining the practicality of boolean tomography with the resolution of analog tomography. In *Proceedings of the 2012 ACM conference on Internet measurement conference*, pages 385–398. ACM, 2012.