Buffering at the Edge: Measuring from Home-routers

Daniel Atkinson, Gustavo Santos, Gaspare Bruno (UFRJ, Brazil)
Edmundo de Souza e Silva (UFRJ, Brazil), Renata Teixeira (INRIA)

1 INTRODUCTION

Video applications have become commonplace in recent years with predictions that, by 2020, approximately 80% of the Internet traffic will be video. Meanwhile, users' expectations concerning the quality of the services they receive have also grown. Questions related to the effect of the home network, the ISP network, and the service providers on the user QoE are of prime concern.

Network measurement efforts have mostly been collected in isolation: at routers in ISP networks, within homes, or at the video service providers. Measurements collected at home routers have mainly been used to understand the quality of service the ISP provides (e.g. [2, 5]), but only a few have addressed the home WiFi [1, 3] and even less addressed the issues related to the combined effect of the home WiFi and the ISP network on the QoE. In particular, understanding the effect of buffer sizes in distinct parts of the network on user QoE is of key importance.

In this work we report on a measurement effort under way in Brazil in partnership with our research group at the university, a midsize ISP and a startup incubated at the university. We focus on a very recent ongoing effort to understand the impact of home-router (HR) buffers on user video QoE and present a simple illustrative example of how the HR WiFi driver buffer size affects the user's perceived QoE.

Initially, we plan to perform laboratory experiments that use models of real time series collected in the ISP network with a monitored video application and consider the impact of different bottlenecks (e.g. wireless, access network). Later we will use ISP volunteer clients.

In Section 2 we describe our measurement efforts in the wild, aiming at providing a sample of part of the analysis we are performing. In Section 3 we show **very preliminary** results from a simple laboratory setup that will serve to illustrate our plans for the near future concerning the impact of home router buffers.

2 MEASUREMENT CAMPAIGN

We partnered with a mid-size ISP in Brazil (Gigalink) and a startup (Anlix) incubated at UFRJ (Brazilian University) to collect active and passive measurements at the HRs of thousands of ISP clients in an experimental joint project with the university. (We do not collect packet headers or any data that could otherwise compromise user's privacy). HRs are the conduit to home devices and, as such, they constitute the ideal place to implement measurement functionalities. The HRs run open software and send information at every minute to servers located in our University laboratory for data analysis. Figure 1 illustrates the measurement infrastructure at each participating residence.

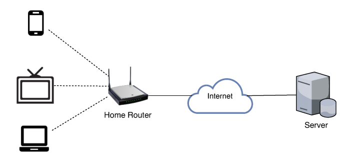

Figure 1: Measurement infrastructure

We collect, at every minute, byte and packet counts from download and upload traffic, round trip packet losses and latency, among others. We have been collecting measurements from approximately 4,000 residences for many months, and plan to collect thousands more in the future as our laboratory infrastructure allows. In addition, we can currently collect basic WiFi metrics (e.g. RSSI, SNR, PHY rate). These constitute several time series of thousands of users that we analyze, to evaluate network performance, user QoE and detect attacks. Moreover, the ISP will provide information concerning its topology and core network routers configuration.

In what follows we briefly mention a few results obtained from our analysis that are related to the buffer studies, in order to provide a sample of our capabilities. Once we create realistic generative models for user traffic we can generate traffic load for the routers in laboratory experiments.

We used unsupervised learning techniques with the objective of discovering users' traffic profiles from download and upload traffic. From the user profiles we can create user traffic generative models. These will be useful in the HR buffer studies we will perform in the laboratory to assess QoE impact.

It is well know that packet losses have a negative impact on user QoE since both startup latency and the number of rebuffering events during video transmission are strongly correlated with losses in the network. Each HR that takes part in the measurement campaign sends, at every minute,

BS'19, December 2-3, Stanford, CA

a burst of 100 ICMP packets at intervals of 10 milliseconds. These ICMP bursts are sent to a server located in the ISP's network. The corresponding time series are analyzed for different purposes. An application, relevant to this work, is to construct models that can emulate real network loss conditions in the laboratory in order to assess HR buffer influence on the QoE at different network loss conditions. The model details are beyond the scope of this work.

3 VERY PRELIMINARY RESULTS

We used a Google Chrome extension, implemented in javascript, by our partner research group, INRIA [4], to collect QoE metrics from youtube (the extension works with many video streaming services but we chose to focus only on youtube for these illustrative experiments). This extension logs javascript video object attributes in JSON format and also collects the HTTP Archive (HAR) file, a log of the browser's HTTP interaction with the website, which is JSON-formatted by default. We then redirected these files to be sent to our local server and created a parser in python to extract relevant QoE metrics, namely, startup delay, mean bitrate, mean resolution, number of changes in resolution and aggregated length of rebuffering events.

For our experiments we chose an action movie trailer, with resolution rates ranging from 144p to 2160p (4K), or 0.085Mbps to 15Mbps (mean bitrate). This video was selected for its wide range of possible resolutions and to try to emulate a "worst case scenario", a very dynamic video where compression capabilities are limited. The desktop where the videos were reproduced was connected to the HR via WiFi. The HR was placed around 6 feet (2 meters) away from the desktop, in an attempt to diminish the impact of other WiFi routers in the vicinity.

In these preliminary experiments, our aim was to simply show that there is good reason to further study the effects of the WiFi buffer size on QoE, so no traffic shaping was performed, only the downstream WiFi buffer size was modified. In the near future, we will perform experiments while varying packet loss rate, latency, bandwidth and cross traffic. We also intend to use the models we developed for the home traffic, packet loss and latency, based on the collected measurements from the ISP clients. Once these experiments have been performed we will have a much better basis for conducting experiments with volunteer ISP customers.

The Figures serve as an illustration of the plans for the investigation. Figure 2 shows the average bitrate of the video decoded at the client as the HR WiFi buffer size varies. Ten experiments were conducted for each buffer size. The blue line is the result of a linear regression fitting to show the trend of the results. As expected, the average bitrate increases with the size of the WiFi buffer and the rate of increase reduces with increasing buffer sizes. One should note that larger buffer sizes contribute to increase the variability of the results.

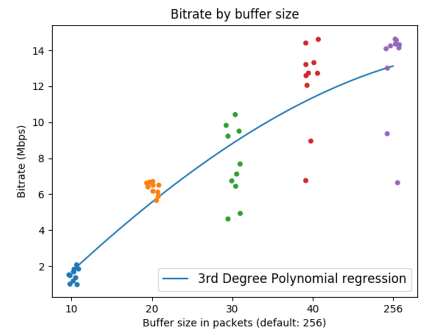

Figure 2: Decoded bit rate vs HR buffer size

Figure 3 shows the startup delay of each of the ten experiments with different HR buffer sizes. As the HR buffer size increases, the startup delay reduces as well as the variability of the delay for each experiment.

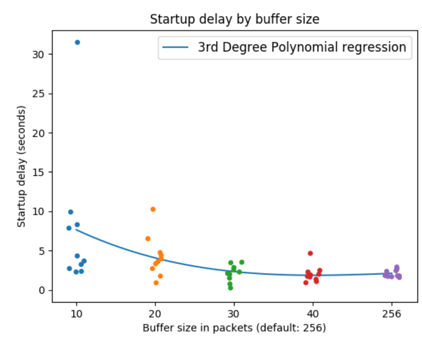

Figure 3: Startup delay vs HR buffer size

Acknowledgements. This work is supported by grants from FAPESP no.2015/24434-5, CAPES, CNPq, FAPERJ.

REFERENCES

[1] D. Da Hora, K. van Doorselaer, Koen Van Oost, and R. Teixeira. Predicting the effect of home Wi-Fi quality on QoE. In INFOCOM 2018.
[2] FCC. Measuring Broadband America.
[3] M. A. Sánchez, J. S. Otto, Z. S. Bischof, D.R. Choffnes, F. Bustamante, B. Krishnamurthy, and W. Willinger. A measurement experimentation platform at the internet's edge. IEEE/ACM Trans. Netw., 23(6):1944–1958, 2015.
[4] P. Schmitt, F. Bronzino, S. Ayoubi, G. Martins, R. Teixeira, and N. Feamster. Inferring streaming video quality from encrypted traffic: Practical models and deployment experience. ACM/POMACS, 2020.
[5] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Measuring home broadband performance. Commun. ACM, 55(11):100–109, 2012.