

QoE Analysis of a Large-Scale Live Video Streaming Event

Adnan Ahmed
The University of Iowa

Zubair Shafiq
The University of Iowa

Amir Khakpour
Verizon Digital Media Services

ABSTRACT

Streaming video has received a lot of attention from industry and academia. In this work, we study the characteristics and challenges associated with large-scale *live* video delivery. Using logs from a commercial Content Delivery Network (CDN), we study live video delivery for a major entertainment event that was streamed by hundreds of thousands of viewers in North America. We analyze Quality-of-Experience (QoE) for the event and note that a significant number of users suffer QoE impairments. As a consequence of QoE impairments, these users exhibit lower engagement metrics.

Keywords

Content Delivery; Live Video; Quality of Experience

1. INTRODUCTION

An increasingly large number of content publishers now broadcast their video content live over the Internet. This growth is a consequence of low costs of content delivery and the adoption of advertisement/subscription based revenue models. Content publishers typically contract third-party Content Delivery Networks (CDNs) to provide high quality over-the-top video streaming services.

According to Cisco [1], more than 64% of the global Internet traffic was video in 2014. The stakeholders in the Internet video ecosystem, such as content providers, CDNs, and ISPs need to frequently upgrade their infrastructure to meet the increasing demands of Internet video and user expectations for high quality. As a result, researchers in academia and industry are actively studying various aspects of over-the-top streaming video delivery.

Prior literature includes studies on analyzing the impact of Quality-of-Service (QoS) metrics, such as bandwidth, packet loss and received signal strength, on the performance of specific applications, including streaming video [11, 6]. More recently, researchers have introduced video-specific Quality-of-Experience (QoE) metrics, such as *rate of buffering* and

average bitrate, in the pursuit to understand and quantify user experience. However, prior studies have been limited mainly to on-demand streaming video. In the recent years, live video streaming has been gaining popularity for various events like sports (MLB World Series, NCAA Basketball Finals etc.[5]) and eSports (Twitch, Azubu etc.[4]). This paradigm, therefore, presents interesting opportunities for video delivery and QoE optimizations due to the unique traffic dynamics and access patterns. Towards this end, we study a large-scale live video streaming event. We focus on a popular entertainment awards event in 2015 that was streamed live by a commercial CDN. We quantify user perceived video quality in terms of QoE metrics such as rate of buffering, average bitrate, rate of fluctuation, etc. We then analyze the statistics of these QoE metrics and observe that a significant number of users suffer from QoE impairments.

2. DATA

Video is by far the biggest source of traffic on the Internet. Video can be classified into three categories: bandwidth-sensitive *streaming stored video* (Netflix, YouTube, Hulu, etc.), bandwidth- and delay-sensitive *streaming live video* (Twitch, NFL Live, CNNgo, etc.), and delay-sensitive *interactive live video* (Skype, Facetime, Google Hangouts, etc.). We focus our attention on *streaming live video* for which while the network must provide the steady-state throughput that is more than the video bitrate, excessive delays can still be an issue due to the live nature of the event. QoE for streaming video is determined by a wide range of factors such as encoding bitrate, buffering rate, etc. Understanding the relationship between QoE and user engagement is an active area of research [3, 9, 2, 10].

Adaptive bitrate controllers are used to dynamically adapt video quality based on various factors such as network bandwidth estimation and player buffer occupancy [7, 12]. The video is encoded at multiple discrete bitrates (quality levels A to G) and divided in segments of 2-10 seconds. The client-side video player estimates the available network bandwidth and requests the video segments at a suitable bitrate. Our data set is composed of server-side HTTP logs from a commercial CDN network which was responsible for online streaming of the live event. The information is collected at multiple Points of Presence (PoPs) across the USA. The HTTP logs contain client information and the requested video chunks. We measure QoE metrics using the timestamps and bitrates of the requested chunks. Overall, our data set contains 625,626 users and more than 21 million minutes worth of video view time. The live event lasted for approximately five hours.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS '16 June 14-18, 2016, Antibes Juan-Les-Pins, France

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4266-7/16/06.

DOI: <http://dx.doi.org/10.1145/2896377.2901504>




Figure 1: Distributions of QoE metrics. Overall, most users experience few buffering events at high average bitrate. We do observe a tail of users experiencing higher buffering, more fluctuations, and lower bitrate than other users.

3. QOE MEASUREMENT

In this section, we conduct a systematic measurement of video QoE and study its impact on user engagement. Commonly used QoS metrics like end-to-end delay, loss and bandwidth, do not directly tie-in to the end-user experience. For example, even if the user is getting low delays at the network level, he/she could be experiencing serious problems at the application level which may arise due to unexpected cross-layer interactions or application layer issues. Therefore, it is important to quantify metrics that directly capture end-user experience. We first define a set of video-specific metrics that are used to quantify QoE [3, 2].

Average Bitrate. Since the video player at the client can request video segments at different bitrates, we use the average bitrate to quantify the overall video quality experienced by the user.

Rate of Fluctuation. We calculate the rate of fluctuation (per minute) as a video QoE metric. Frequent changes in video bitrate may degrade user experience.

Rate of Buffering. We calculate the rate of buffering as the number of buffering events per minute over the duration of the video session. As reported in prior literature [3, 9, 8], buffering events can have a significant negative impact on user engagement.

Buffering Ratio. We calculate buffering ratio as the fraction of time spent in buffering events over the duration of the video session.

Figure 1 plots the cumulative distribution functions (CDFs) of the QoE metrics. In Figure 1(a), we note that a vast majority of users experience relatively low rate of buffering. For instance, almost 95% of users experience less than 0.5 buffering events per minute on average. Similarly, in Figure 1(b), we note low buffering ratios for most users. More specifically, more than 90% of users have buffering ratio lower than 0.05. In Figure 1(c), we observe that users do experience relatively more video quality fluctuations than buffering events. For example, about 50% of users experience more than 2 fluctuations per minute on average. In Figure 1(d), we note that a majority of users watch the video stream at a high bitrate. In particular, more than 60% of users watch the video stream at a bitrate of at least 1 Mbps (i.e., quality level F) on average. Overall, we observe that a small number of users in the tail of QoE distributions suffer from frequent buffering events, frequent video quality fluctuations, and low average bitrate. Our further analysis shows that these “tail” users have lower engagement than other users.

4. CONCLUSION

In this paper, we analyzed QoE metrics for a large-scale live video streaming event. We analyze server side logs to extract video QoE metrics for more than 600K users across USA for a large-scale live video streaming event. We note that a significant number of tail users suffer QoE impairments. As reported in prior literature, we observe that these tail users have lower engagement compared to other users.

Acknowledgement

This material is based in part upon work supported by the National Science Foundation under Grant Number CNS-1464110.

5. REFERENCES

- [1] Cisco Visual Networking Index: Forecast and Methodology, 2014-2019.
- [2] A. Balachandran, V. Sekar, A. Akella, S. Seshan, I. Stoica, and H. Zhang. Developing a Predictive Model of Quality of Experience for Internet Video. In *ACM SIGCOMM*, 2013.
- [3] F. Dobrian, A. Awan, D. Joseph, A. Ganjam, J. Zhan, V. Sekar, I. Stoica, and H. Zhang. Understanding the Impact of Video Quality on User Engagement. In *ACM SIGCOMM*, 2011.
- [4] L. Eadicicco. 10 Facts About Twitch. <http://www.businessinsider.com/statistics-about-twitch-2014-8>, August 2014.
- [5] M. Fisher. The Next Chapter in Live Streaming. <http://qwilt.com/the-next-chapter-in-live-streaming-in-which-certai>, April 2014.
- [6] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. YouTube Traffic Characterization: A View From the Edge. In *ACM Internet Measurement Conference (IMC)*, 2007.
- [7] T.-Y. Huang, R. Johari, N. McKeown, M. Trunnell, and M. Watson. A Buffer-Based Approach to Rate Adaptation: Evidence from a Large Video Streaming Service. In *ACM SIGCOMM*, 2014.
- [8] J. Jiang, V. Sekar, I. Stoica, and H. Zhang. Shedding Light on the Structure of Internet Video Quality Problems in the Wild. In *ACM Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, 2013.
- [9] S. S. Krishnan and R. K. Sitaraman. Video Stream Quality Impact Viewer Behavior: Inferring Causality Using Quasi-Experimental Designs. In *ACM Internet Measurement Conference (IMC)*, 2012.
- [10] M. Z. Shafiq, J. Erman, L. Ji, A. X. Liu, J. Pang, and J. Wang. Understanding the Impact of Network Dynamics on Mobile Video User Engagement. In *ACM SIGMETRICS*, 2014.
- [11] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, and J. Wang. A First Look at Cellular Network Performance during Crowded Events. In *ACM SIGMETRICS*, 2013.
- [12] X. Yin, A. Jindal, V. Sekar, and B. Sinopoli. A Control-Theoretic Approach for Dynamic Adaptive Video Streaming over HTTP. In *ACM SIGCOMM*, 2015.