

Modeling the Effect of Sharing Policies for Network-assisted HTTP Adaptive Video Streaming

Jan Willem Kleinrouweler*, Sergio Cabrero*, Rob van der Mei†*, Pablo Cesar*

*Centrum Wiskunde & Informatica
Science Park 123
1089 XG Amsterdam,
The Netherlands

†VU University
De Boelelaan 1081
1081 HV Amsterdam,
The Netherlands

{j.w.m.kleinrouweler, cabrero, r.d.van.der.mei, p.s.cesar}@cwi.nl

1. INTRODUCTION

HTTP adaptive streaming (HAS) is becoming the most popular paradigm for video streaming over the Internet. In HAS a video file is split-up into segments that are encoded at multiple bitrates. The segments are published on a Web server, alongside a manifest file that describes in what bitrates the video is available. Depending on the bandwidth that is available, a video player selects a matching bitrate. This approach gained much interest from both research and industry, because it allows to do adaptive streaming with the appealing properties of streaming over HTTP: firewall traversal and usage of off-the-shelf Web servers.

Despite its popularity, it was found that selecting an appropriate bitrate is a difficult task for a player [4], especially when multiple players share a bottleneck link [1]. As a response and acknowledging *Quality of Experience* (QoE) as an important factor, a plentitude of solutions has been proposed in the literature in the last few years. This also gave rise to a group of solutions that can be classified as *network-assisted HAS* [3, 6, 2, 5]. These solutions employ in-network devices and use their knowledge about traffic over the bottleneck link to assist HAS players in selecting a bitrate. The HAS-controllers are typically located at one of the edges of a bottleneck link, for example in the residential gateway of a home network.

The solutions that use in-network information to assist HAS players show promising results and make it interesting to further study this approach. It opens up possibilities to create managed video delivery environments where players are treated differently based on their characteristics, creating an optimized viewing experience for all users in the network. The streams' bitrates and variations in these bitrates have been identified to play a major role in the QoE of the viewer [7]. However, building testbeds and evaluating the performance of sharing policies is a costly and time consuming process. Therefore, we proposed a Markov model to determine the mean bitrate and bitrate variations of HAS players, and successfully applied this model in a configuration where the available network capacity is equally divided among all players [5].

In this paper we continue this work and present how our model can also be used for analyzing other sharing policies that can be executed by a HAS-controller. By means of an example that shows how devices with different screen sizes get different bitrates assigned, we demonstrate that the model based results are highly accurate when comparing them to a real implementation. As such, our model is a useful tool in planning video delivery networks and development of sharing policies.

2. PERFORMANCE MODEL

Starting and stopping video players change the demand on the shared network connection, and it may requires the other players to change the quality of their streams to fit all active players. A HAS-controller takes the number of players, the types of devices and the stream characteristics into account to divide the available capacity among the players, according to a certain policy. Every time a new player starts or an existing player stops the capacity is redivided. In [5] we propose a Markov model for this process that outputs the mean bitrate and how often the bitrate changes for managed HTTP adaptive streams. In our model, players that are treated the same according to a policy are grouped. Each state in the Markov process represents a combination of players and is defined as a vector (n_1, n_2, \dots, N_k) , where n_k is the number of players of group k , and K the number of different groups that is considered in a policy. The state space is limited to all states with non-negative integer-valued entries that satisfy the following condition:

$$\sum_{k=1}^K n_k \bar{B}_k \leq C,$$

where C is the capacity of the shared link and \bar{B}_k the lowest available bitrate for players in group k . For each state, and each group k , a bitrate is selected by a policy function based on the number of players, available capacity, and available video bitrates. Changing the policy in the model is straightforward, and only requires to modify the function that assigns a bitrate to each group in each state. From the fraction of time spend in each state and the bitrates assigned to each group of players, we can obtain the mean bitrate. The number of transitions between states that have different bitrates assigned to a group determine the number of bitrate switches.

3. DEVICE HETEROGENEITY

For devices with different screen sizes it is not fair to equally divide the available capacity among the players. Reference [2] describes how different bitrates can be compared among devices with different form factors to yield an equal viewing experience. In our analysis we use the same device classes and a similar bitrate mapping between the classes. The first group represents smartphone sized devices that stream a 360p video of 60 seconds encoded at 400, 600, and 1000 kbit/s. The second group represents tablet viewers that stream a 720p video of 120 seconds encoded at 400, 600, 1000, 1500, 2000 kbit/s. The third group is large screen devices that stream a 1080p video of 180 seconds encoded at 400, 600, 1000, 1500, 2000, and 4000 kbit/s.

Figure 1 shows the model based comparison between a policy that equally divides the bandwidth among the players (Policy1) and




Figure 1: Model-based comparison of a non device aware and a device aware policy




Figure 2: Comparison of the model-based performance with the actual performance achieved through experiments for the device aware policy

a policy that takes the form factors of the devices into account (Policy2). The results clearly show that under Policy2 the small screen devices make room for the large screen devices, and the number of switches in quality is lowered, except for small screen devices at low arrival rates. At lower arrival rates the small screen devices have to switch to a lower bitrate first, so that large screen devices can remain streaming at higher bitrates for a longer time. At higher arrival rates the 360p devices are most of the time already at the lowest bitrate, and it requires the larger screen devices to make quality switches.

Figure 2 compares the model-based performance with the actual performance that we achieved in experiments using our proxy-server based HAS controller that implements Policy2 (testbed details in [5]). The results show that our model is highly accurate for both mean bitrate and the number of quality switches.

4. CONCLUSION

Video streaming over the Internet is becoming so popular that it is no longer an exception that users of a shared network connection stream videos at the same time. If this network connection has limited capacity, it is important think about how this capacity can be shared to provide an optimal viewing experience to all users. With the example in this paper we demonstrate that our model is an accurate tool that can be used to analyze the performance of sharing policies in managed HTTP adaptive streaming environments. The model is intended to be used in dimensioning video delivery networks and in the development of sharing policies.

5. REFERENCES

- [1] S. Akhshabi, L. Anantakrishnan, A. C. Begen, and C. Dovrolis. What happens when HTTP adaptive streaming players compete for bandwidth? In *NOSSDAV '12: Proceedings of the 22nd international workshop on Network and Operating System Support for Digital Audio and Video*, pages 9–14, New York, New York, USA, June 2012. ACM Request Permissions.
- [2] P. Georgopoulos, Y. Elkhatib, M. Broadbent, M. Mu, and N. Race. Towards network-wide QoE fairness using openflow-assisted adaptive video streaming. In *FhMN '13: Proceedings of the 2013 ACM SIGCOMM workshop on Future human-centric multimedia networking*, pages 15–20, New York, New York, USA, Aug. 2013. ACM Request Permissions.
- [3] R. Houdaille and S. Gouache. Shaping HTTP adaptive streams for a better user experience. In *MMSys '12: Proceedings of the 3rd Multimedia Systems Conference*, pages 1–9, New York, USA, Feb. 2012. ACM Request Permissions.
- [4] T.-Y. Huang, N. Handigol, B. Heller, N. McKeown, and R. Johari. Confused, timid, and unstable: picking a video streaming rate is hard. In *IMC '12: Proceedings of the 2012 ACM conference on Internet measurement conference*, pages 225–238, New York, New York, USA, Nov. 2012. ACM Request Permissions.
- [5] J. W. Kleinrouweler, S. Cabrero, R. van der Mei, and P. Cesar. Modeling Stability and Bitrate of Network-Assisted HTTP Adaptive Streaming Players. In *27th International Teletraffic Congress (ITC 27)*, Ghent, Belgium, Sept. 2015.
- [6] R. K. P. Mok, X. Luo, E. W. W. Chan, and R. K. C. Chang. QDASH: a QoE-aware DASH system. In *MMSys '12: Proceedings of the 3rd Multimedia Systems Conference*, pages 11–22, New York, New York, USA, Feb. 2012. ACM Request Permissions.
- [7] M. Seufert, S. Egger, M. Slanina, T. Zinner, T. Hossfeld, and P. Tran-Gia. A Survey on Quality of Experience of HTTP Adaptive Streaming. *Communications Surveys Tutorials, IEEE*, PP(99):1–1, 2014.