Client Starvation: A Shortcoming of Client-driven Adaptive Streaming in Named Data Networking

Daniel Posch, Christian Kreuzberger, Benjamin Rainer and Hermann Hellwagner
Institute of Information Technology (ITEC)
Alpen-Adria-Universität (AAU) Klagenfurt, Austria
firstname.lastname@itec.aau.at

ABSTRACT

Information-centric Networking (ICN) as a potential Future Internet architecture has to efficiently support the consumption of multimedia content. Recent proposals consider the reuse of MPEG-DASH to provide adaptive streaming in ICN. Due to the fact that MPEG-DASH relies on pure client-driven adaptation, it encounters difficulties dealing with ICN's inherent caching and multi-path transmission. By conducting simulations using the concrete ICN approach Named Data Networking (NDN), we show that pure client-driven adaptation leads to shortcomings. Furthermore, we propose to use in-network adaptation based on scalable content for overcoming these shortcomings in NDN.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Network communications

General Terms

Design, Theory

Keywords

Information-centric Networking; Adaptive Streaming

1. INTRODUCTION

In today's Internet, multimedia content represents the biggest traffic source. This draws the attention of ICN research towards the question of how to realize effective audiovisual content dissemination. Due to its effectiveness in IP-based networks, the integration of MPEG-DASH (ISO/IEC 23009-1) in NDN has been investigated by recent research [3]. MPEG-DASH is able to consider consumer demands, e.g., heterogeneous end-devices, and responds well to fluctuating network conditions. It uses an elegant approach: Content is split into independent segments, which are encoded in various quality levels. Complexity is pushed to the clients, which are responsible for selecting and fetching appropriate segments. The adaptation process considers only local parameters. ICN's specifics, e.g, network inherent caching and multipath delivery, are disregarded. This causes negative side-effects on other network participants. One example is *client starvation*, which is presented in this article. It describes the fact that requests for popular content overwhelm requests for less popular content, providing insufficient service for some clients.

ICN'14, September 24–26, 2014, Paris, France.
ACM 978-1-4503-3206-4/14/09.
http://dx.doi.org/10.1145/2660129.2660162.

2. EXPERIMENTAL SETUP

The objective of adaptive streaming is the enhancement of user-perceived quality, also referred to as Quality of Experience (QoE). As we consider only reliable ICN transport, the most influential QoE factors are [6]: (1) startup delay, (2) playout stalling time, (3) number of quality switches and (4) playback quality.

Using ns-3/ndnSIM [1], we set up a simulation with 100 MPEG-DASH clients, streaming 10 different videos over a bottleneck link (20, 25, 30 Mbit/s). The node was equipped with a 256 MB LRU content store. For each evaluated bottleneck-speed we perform 30 simulation runs and for each of the runs the video streamed by a client is randomly assigned based on the popularity distribution depicted in Table 1. Clients start streaming at varying times, following an exponential distribution.

The videos are encoded according to the Scalable Video Coding (SVC) extension of H.264/MPEG-4 AVC [4] using SNR-scalability. The encoder was configured to provide contents in six different quality layers, resulting in the following average bitrates: 1290, 1768, 2220, 2608, 2933 and 3197 kbit/s. We selected SVC-based encoding because ICN caches can be used more efficiently compared to non-scalable content encodings.

The clients use a buffer-based adaptation logic as described in [5]. Additionally a TCP-like congestion-control window (AIMD) was used to adapt a client's data request rate. We justify this step by the latest developments in the Future Internet community [2], which argues for a dual-stack deployment of ICN and TCP/IP.

Content consumed	by X Clients
Video 0 / Video 1	1 / 1
Video 2 / Video 3 / Video 4	2 / 2 / 2
Video 5 / Video 6 / Video 7	4 / 8 / 16
Video 8 / Video 9	32 / 33

Table 1: Popularity distribution of the videos.

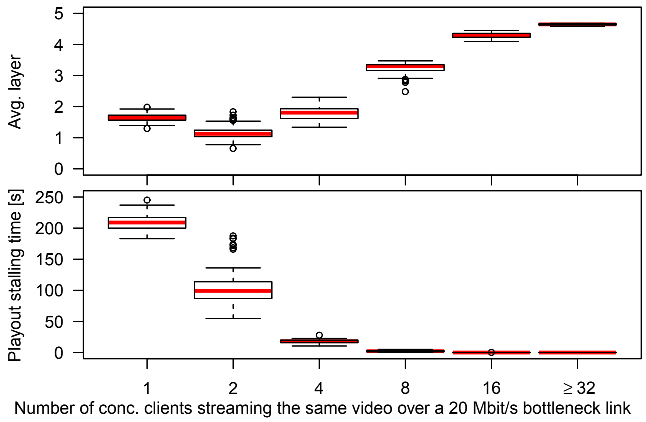

Figure 1: Average received layer and stalling time versus video popularity.

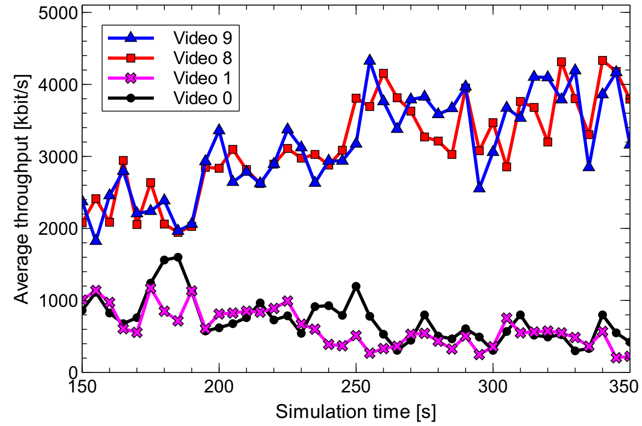

Figure 2: Average bandwidth share on the bottleneck link per video for a single simulation run.

3. RESULTS

Figure 1 depicts the results of the simulation with a bottleneck of 20 Mbit/s via box plots (we omit similar results for 25 and 30 Mbit/s due to space constraints). The X-axis represents the number of concurrent clients according to Table 1, while the Y-axis shows the playout stalling time and the average received layer. The box plots show that the more popular a video is, the better the received quality and the lower the number of stalls on average is. Exceptions are videos requested by only one client. They have a higher average quality than videos requested by two clients. Due to the high playout stalling time, videos requested by only one client receive more bandwidth towards the end of the simulation.

The higher quality levels of clients requesting popular videos can be partly explained by efficient cache utilization in ICN, however, the long playout stalling duration of some clients has another cause. In the case of congestion, ICN nodes discard packets due to buffer overflows. Therefore, requests stay unsatisfied, resulting in timeouts, and clients reduce their congestion window to avoid future congestion. However, if multiple clients request the same content, dropped requests are likely to be re-issued by other clients. This may result in delivery of data even if a previous request has been discarded. Consumers who stream the same video as many other clients, may not notice congestion. Therefore they will not reduce their congestion window, and continue to utilize a higher bandwidth share. Furthermore, cache hits increase the congestion window's size, and once cache misses are encountered, those clients demand a too large share of the available bandwidth. The simulations indicate that clients requesting unpopular content starve, as illustrated by the high stalling duration in Figure 1. We call this problem *client starvation*. In addition, Figure 2 shows that the average bandwidth share per video on the bottleneck link is higher for popular than for unpopular videos. For illustration purposes only a subset of the videos are depicted during 200 seconds of simulation time.

4. CONCLUSION AND FUTURE WORK

We outlined the problem of *client starvation*. This issue is mainly caused by the fact that the clients are not aware of how many other concurrent users are consuming the same video, nor do they know anything about the received video quality of other clients. Therefore client-driven adaptation can not solve *client starvation* satisfactorily. As this problem emerges within the network, we propose the usage of in-network adaptation (INA) to tackle the problem at its place of origin, which is part of our future research. We are going to investigate INA in ICN based on layered content encodings. This enables INA despite ICN's content-based security model, which prohibits content manipulation due to security constraints introduced by digital signatures.

5. ACKNOWLEDGMENTS

This work was partially funded by the Austrian Science Fund (FWF) under the CHIST-ERA project CONCERT (A Context-Adaptive Content Ecosystem Under Uncertainty), project number *I1402*.

6. REFERENCES

[1] A. Afanasyev, I. Moiseenko, and L. Zhang. ndnSIM: NDN Simulator for NS-3. Tech. Report NDN-005, 2012.

[2] S. Braun, M. Monti, M. Sifalakis, and C. Tschudin. CCN & TCP Co-Existence in the Future Internet: Should CCN be Compatible to TCP? In *IFIP/IEEE International Symposium on Integrated Network Management*, 2013.

[3] S. Lederer, C. Müller, B. Rainer, C. Timmerer, and H. Hellwagner. Adaptive Streaming over Content Centric Networks in Mobile Networks using Multiple Links. In *Proceedings of the IEEE International Conference on Communication (ICC)*, 2013.

[4] H. Schwarz, D. Marpe, and T. Wiegand. Overview of the Scalable Video Coding Extension of the H. 264/AVC Standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 17(9):1103–1120, 2007.

[5] C. Sieber, T. Hoßfeld, T. Zinner, P. Tran-Gia, and C. Timmerer. Implementation and User-centric Comparison of a Novel Adaptation Logic for DASH with SVC. In *IFIP/IEEE Workshop QCMan*, 2013.

[6] L. Yitong, S. Yun, M. Yinian, L. Jing, L. Qi, and Y. Dacheng. A Study on Quality of Experience for Adaptive Streaming Service. In *IEEE International Conference on Communication (ICC)*, June 2013.