Keep Calm and Don't Switch: About the Relationship Between Switches and Quality in HAS

Christian Moldovan, Tobias Hoßfeld
Chair for Modelling of Adaptive Systems,
University of Duisburg-Essen, Germany
{christian.moldovan, tobias.hossfeld}@uni-due.de

Korbinian Hagn, Christian Sieber, Wolfgang Kellerer
Chair of Communication Networks,
Technical University of Munich, Germany
{k.hagn, c.sieber, wolfgang.kellerer}@tum.de

***Abstract*—Video streaming has established itself as the main method to consume multimedia content on the Internet. The most widespread technique is HTTP Adaptive Streaming which is used by large video service platforms such as YouTube and Netflix. In order to guarantee a high QoE, different adaptation strategies have been developed that lead to different video playout patterns. While some strategies may be more aggressive than others and often adapt the video quality, it is not clear which strategy is the best.**

In this paper, we want to identify the trade-off between the average video quality and switches in the quality during playout. We do this with a user-centric view and try to optimize the adaptation depending on the user preference with a quadratic program. This work allows us to put existing and future video adaptation algorithms in perspective with respect to user preferences. Our results show that the video quality can already be increased greatly by allowing few switches while more switches lead to diminishing gains. This is a novel discovery that is important for user-centric QoE-management which is of high interest for ISPs and video service providers.

***Index Terms*—Video streaming, quality, switches, optimization, quadratic programming**

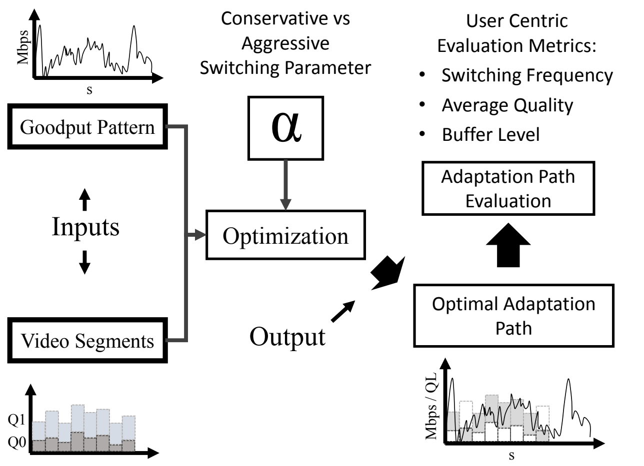

Figure 1. Contribution and methodology of this paper. We propose an optimization formulation which calculates an optimal HTTP adaptive streaming segment picking strategy (*adaptation path*) for a given goodput pattern and a switching aggressiveness parameter α. We evaluate the resulting adaptation paths based on metrics related to the user experience and based on the selected value for α.

I. Introduction

In the last few years a trend became visible, to go from traditional television to Internet based video services. With video streaming becoming the status quo in media consumption, quality expectations are increasing. Low quality videos are no longer considered acceptable in contrast to some years ago. Therefore, Internet Service Providers (ISPs) and video service providers are facing the challenge of providing seamless multimedia delivery in high quality.

In adaptive video streaming the quality in which the next segment is downloaded is determined by a number of factors such as the state of the buffer or the current throughput. This may lead to a high number of quality switches, depending on the adaptation strategy. Even though stalling can be avoided in most cases, a high number of quality switches may also have a negative impact on the Quality of Experience (QoE) [1], [2]. A very passive adaptation strategy plays the video on a lower quality in order to keep the number of switches at a minimum. While some users may prefer watching a video with a high average quality with many switches over low quality with few switches, there is little research on the trade-off between quality and switches and user preferences.

In this paper we analyze the trade-off between average video quality and quality switches with regard to a weighting parameter. For this purpose, we use a quadratic program, that optimizes adaptation in video streaming towards higher quality and fewer quality switches. This is done with respect for the users individual preference for these two QoE impact factors. Our research question can be formulated as follows. *What is the trade-of between the average quality of a video and the number of quality switches in adaptive video streaming?*

Figure 1 summarizes the contribution and methodology of this paper. We demonstrate the optimization of the adaptation with a real mobile goodput trace and 41 different YouTube videos. Various user preferences for the adaptation are respected by conducting a parameter study for the adaptation aggressiveness parameter α that defines how frequently the player may switch to another quality. We then investigate the resulting optimal adaptation paths and evaluate key QoE indicators, such as the switching frequency, the average video quality and the buffer level.

The next section discusses background and related work. In Section III, the optimization problem is proposed. Thereafter, the methodology of our evaluation is presented in Section IV. In Section V, we present the results of our evaluation and discuss them in detail. The following section gives a summary of the results of the paper and discusses the implications of the findings. Finally, we conclude the paper and give an outlook on future work.

II. Background & Related Work

In HTTP adaptive streaming, a video is divided into segments of equal duration. Each segment is available in different bit rates that result in different quality. Each time a new segment is requested, based on the adaptation strategy it is selected on which quality the segment is downloaded. This choice is usually determined by a heuristic that relies on the buffer state [3] or bandwidth estimations [4], [5]. Such adaptation heuristics differ to a large degree in their implementation. Some are aggressive and try to switch to a higher quality early while others try to avoid many quality switches. A survey of key rate adaptation techniques is presented in [6]. An overview of the most important HAS algorithms is given in [7]. A detailed analysis of the adaptation of YouTube is given in [8].

Reference [2] investigates the impact of quality switches on the QoE in comparison to stalling events. Their results show that smooth switching only performs slightly better than abrupt switching. They find that stalling and adaptation had similar impact on the QoE. Furthermore, it is discovered that a high number of quality switches does not lead to a significantly lower QoE. In contrast, we only focus on application-layer Quality of Service (QoS), as we investigate to what degree the video quality can be increased if a high switching frequency is allowed. Comprehensive surveys of important papers on QoE are conducted in [9], [10]. Among others, the authors discuss the impact of layer switches on the QoE. They consider switching itself as a degradation and agree with other authors [1], [11] that the number and amplitude of switching events should be kept low. Furthermore, they find that the time that is spend on each quality layer also impacts the QoE. In [12] the impact of the switches and of the time on the highest layer on the QoE is compared. The authors find that the time spent on the highest layer is identified as the main influence factor while the switches have no significant impact. As a conclusion, they omit the switches from their QoE model. This dissent concerning the importance of quality switches for the QoE shows us that research in this area is not fully explored and is still progressing. For this reason, we want to investigate at what cost in terms of quality we can decrease the number of switching events. What cost users are willing to pay is subject of future work and is considered an unknown parameter in this work.

In order to optimize the adaptation in video streaming, we use a quadratic program which has already been presented in previous work: [13] and [14] investigate a mobile video streaming scenario in which a mobile client enters a tunnel. Even though no data can be downloaded in the tunnel, no stalling must occur in the video. To solve this problem optimally, a quadratic program was defined which is compared to an adaptation heuristic. The authors of [15] discuss the same problem, but use a two-step approach for their optimization. They propose to first maximize the mean quality and then minimize the number of switches if possible. Furthermore, a special linear case of the program was used in [16]. In this previous work, the authors investigate how much data is downloaded redundantly by YouTube's adaptation heuristic. The application-layer QoS of this scenario is compared to the QoS in an optimal scenario with no redundantly downloaded data. For this second scenario, the aforementioned linear program was used. Another linear program that tries to solve the problem of optimal rate allocation for video streaming in mobile networks is discussed in [17].

III. Problem Formulation

In the following, we present an exact formulation of the problem that we want to optimize in this paper. Consider a video that consists of n segments. Each segment i is downloaded in exactly one of r quality layers. In order to play a video without stalling, each segment i must be downloaded before its deadline D_i. The data that is available from the initial video request at a point in time t is defined as $V(t)$. The initial startup delay is fixed to 5 s in the evaluation, i.e. $V(0)$ equals the sum of the goodput of the first 5 s. The size of segment i on layer j is S_{ij} and is given in Byte. If two consecutive segments are downloaded on different layers, the viewer experiences a quality switch. In order to maximize the viewers quality of experience, we want to increase the mean quality and reduce the number of quality switches. The importance of these two parameters is set as $\alpha \in [0, 1]$. Higher values for α indicate that it is more important to avoid quality switches than to increase the average quality. Different users have different preferences in this regard and thus different values for α. If a segment i is downloaded on layer j, then we define $x_{ij} := 1$, otherwise $x_{ij} := 0$.

The goal is to decide, on which quality layer we must download each segment in order to maximize the weighted sum of mean quality and the number of quality switches while avoiding stalling. This optimization problem can be formulated as a quadratic program as follows.

$$\max \quad \sum_{j=1}^{r} \left(\frac{\alpha}{nr} \sum_{i=1}^{n} j x_{ij} - \frac{1-\alpha}{2(n-1)} \sum_{i=1}^{n-1} (x_{ij} - x_{i+1,j})^2 \right) \tag{1}$$

$$\text{s.t.} \quad \sum_{j=1}^{r} x_{ij} = 1 \;\; \forall i \in \{1, \dots, n\} \tag{2}$$

$$x_{ij} \in \{0, 1\} \;\; \forall i \in \{1, \dots, n\},\; j \in \{1, \dots, r\} \tag{3}$$

$$\sum_{i=1}^{k} \sum_{j=1}^{r} S_{ij} x_{ij} \leq V(D_k) \;\; \forall k \in \{1, \dots, n\} \tag{4}$$

The objective function (Equation 1) maximizes the weighted sum of the mean quality and the number of quality switches. In order to receive values between 0 and 1, we normalize the mean quality by the maximum quality r, we normalize the

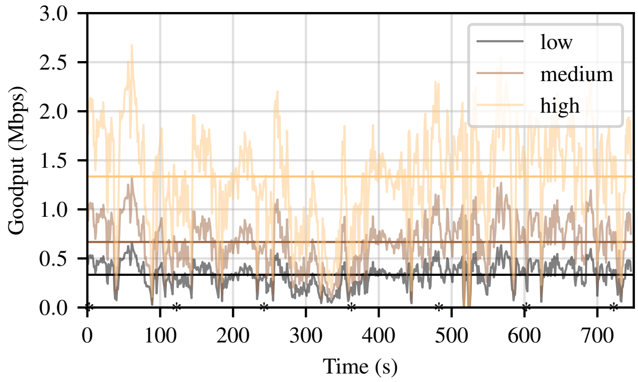

Figure 2. Low, medium and high goodput traces from a vehicular mobility scenario with means of 0.33 Mbps, 0.67 Mbps and 1.34 Mbps. The horizontal lines show the mean goodput of the three traces. The coefficient of variation (CV) is the same for the three traces. The stars at the bottom mark the different starting timestamps of the video playback used in the evaluation.

switches by the highest possible number of switches $n-1$ and we add the factor $1/2$ to the quadratic term since it increases by 2 with every switch. Constraint 2 and 3 ensure that each segment is download in exactly one quality. Constraint 4 ensures that each segment k is downloaded before its deadline D_k while not more data than $V(D_k)$ is downloaded.

IV. Evaluation Methodology

The primary objective of the evaluation is to assess the influence of the parameter α on the average playback quality, on the number of switches and on the buffer level during the playback. For the evaluation we choose a challenging scenario based on a mobile goodput trace collected while riding by car through a city. The segment sizes, i.e. the videos, were selected from YouTube based on specified properties, e.g. based on duration and number of quality levels. In the following we first introduce the mobile pattern in detail. Afterwards we discuss the selection of the videos.

A. Mobile Goodput Trace

The goodput trace defines the instantaneous application-layer downlink throughput of a client's network connection. The trace was collected in December 2012 while driving by car in and around Klagenfurt, Austria. A single large HTTP GET request was send via a UMTS stick and the resulting goodput of the download recorded continuously for a duration of 750 s. For the evaluation, the original recorded trace is scaled to a mean of 0.33 Mbps, 0.67 Mbps and 1.34 Mbps, while keeping the coefficient of variation (CV) the same (0.38). The patterns have an autocorrelation of 0.80 for a lag of 1. We denote the resulting three patterns as *low*, *medium* and *high*. Furthermore, we define seven shifted versions of the patterns where we move the starting point forward and append the skipped goodput samples to the end of the patterns. The starting timestamps are $\{0\,\mathrm{s},\ 120\,\mathrm{s},\ ...,\ 720\,\mathrm{s}\}$. Thus, we use in total $7 \cdot 3 = 21$ goodput patterns in the evaluation for each video sequence.

Figure 2 depicts the three patterns. The average goodput of each pattern is indicated as horizontal bar. The stars at the lower axes indicate the three shifting timestamps. The figure shows the challenging nature of the vehicular mobility patterns in terms of goodput variation. The goodput drops frequently to less than 0.1 Mbps and two times, at about 510 s into the measurement, to zero. Furthermore, there is no period of more than 60 s without a significant drop. The maximum observed goodput per pattern is 0.67 Mbps, 1.34 Mbps and 2.68 Mbps. The traces of the three patterns are available in the supplemental material to the paper [18].

B. Videos

We use in total 41 different videos for the evaluation. The videos represent the content mix of the videos uploaded to YouTube and are randomly chosen by popularity at time of the study. The characteristics of the video set is summarized in Table I. The videos have a length of one to ten minutes (average 5.3 minutes) and are from different video categories (*"minecraft", "music", "funny cats", "gopro", "game"*). All videos have five quality levels ($\{144p, 240p, 360p, 480p, 720p\}$) with an average bitrate ranging from about 0.1 Mbps for the lowest quality level to 1.3 Mbps for the highest quality level. Based on the downloaded video files, we split the videos in segments with a duration of 5 seconds and use the segment sizes as input for the optimization. For details about the video selection process we refer the reader to a previous study [19] where the same videos were used.

Table I
Video Testset Characteristics

Avg. bit-rates	140p	240p	360p	480p	720p
Min (Mbps)	0.08	0.16	0.04	0.08	0.14
Mean (Mbps)	0.10	0.23	0.36	0.68	1.33
Max (Mbps)	0.11	0.24	0.56	1.05	2.08
Std	0.00	0.01	0.13	0.24	0.48
Segmentsize	5s				
Duration	{1, 2, ..., 10} minutes, avg 5.3 m				

Figure 3 illustrates the progression of the bit-rate through the video as an example for the video *CRZbG73SX3s*, a first-person-view sports clip. The average goodput of the three traces from Figure 2 are shown as transparent horizontal lines. The figure shows that the maximum bit-rate is limited by the encoder. Furthermore the figure depicts that low motion or low detail scenes are encoded more efficiently and the average bit-rate drops frequently. The two lower quality levels do not exhibit any significant variations in bit-rate compared to the upper three levels. The video segment sizes are available in the supplemental material to the paper.

V. Evaluation

Under the assumption that the player does not do breaks in between segment downloads and that the player buffer is unlimited, the following holds true. An aggressive switching behavior results in a high average quality, high number of switches and a low average buffer level. A conservative

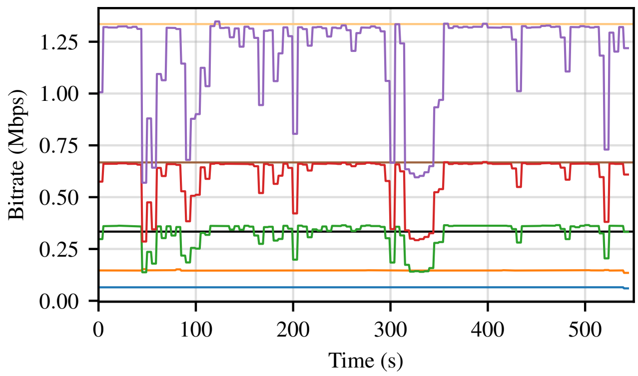

Figure 3. Example video segments for video CRZbG73SX3s and the quality levels $\{144p, 240p, 360p, 480p, 720p\}$ (from bottom to top). Each segment has a duration of 5 s. The horizontal bars mark the average goodput of the traces in Figure 2. The two lowest levels do not exhibit significant variations, the three highest levels exhibit bit-rate drops in low motion/detail scenes.

switching behavior decreases the average quality, decreases the number of switches and increases the average buffer level in the player. From this it follows, that the main objectives of a QoE-aware streaming player, i.e. to increase the average quality, to decrease the number of switches and to avoid stalling by keeping the buffer level high, are contradictory. The main question of the evaluation is: *Can we keep the average quality high while at the same time reduce the number of switches and increase the average buffer level?*

In the following we first discuss the influence of the α parameter by example. Figure 4 illustrates the adaptation path for three different values of α for the video CRZbG73SX3s under the medium traffic pattern and a starting timestamp of 0 s. $\oslash$ denotes the average quality of the playback. It can be observed that for $\alpha = 0$ the adaptation path is very conservative as there are zero quality switches and quality level 3 is selected from start until the end of the playback. For $\alpha = 0.5$, the number of switches increases to five and the average quality to 3.29 as the adaptation path is able to show

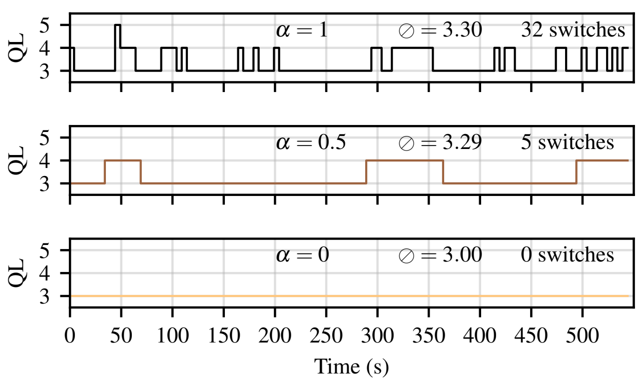

Figure 4. Adaptation path for three different values of α for the example video CRZbG73SX3s, medium goodput pattern and a starting timestamp of 0 s. $\oslash$ denotes the average quality. An increase in α increases the switches and the quality. From $\alpha = 0.5$ to $\alpha = 1.0$ the increase in quality is only 0.01 quality levels, while the switching frequency increases dramatically.

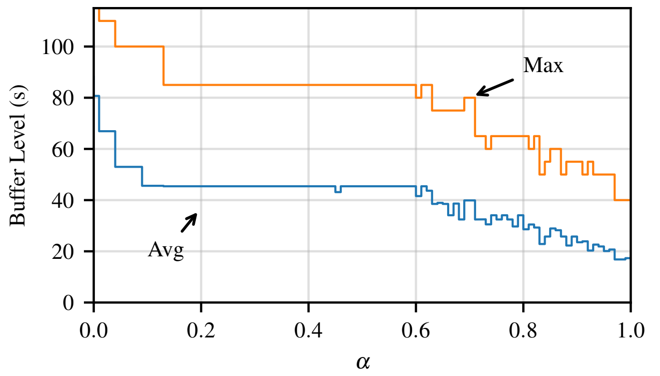

Figure 5. Average and maximum buffer level for video CRZbG73SX3s for the medium goodput pattern. Aggressive choices of α decreases the observed average buffer level compared to very conservative choices. Results are presented as median over all seven starting timestamps.

for three periods during the playback a higher quality level. For $\alpha = 1$, the adaptation path is very aggressive with 32 quality switches, i.e. about 3.5 per minute, and one period where even quality level 5 is selected. However, the average quality increase is marginal with 0.01 compared to five switches for $\alpha = 0.5$. It follows that in this example the α parameter is able to adjust the aggressiveness of the adaptation path. Furthermore, we see that a high number of switches is not necessarily helpful in increasing the average quality.

A. *Observations for example video CRZbG73SX3s*

Next we illustrate the relationship between the choice of α and the average and maximum buffer level by example. Figure 5 shows the average and maximum buffer level in seconds for video CRZbG73SX3s for the medium goodput pattern. The average buffer level is the time-dependent average over the buffer level values observed during the playback. The maximum is the highest buffer level observed during playback.

Two major observations can be made from the figure. First, the buffer level decreases for more aggressive values of α. For a conservative choice of alpha, e.g. $\alpha = 0.05$, the buffer level is on average about 53 s and maximum 100 s. For an aggressive choice, e.g. $\alpha = 1.0$, the buffer level is only 17 s on average and a maximum of 40 s is observed. The second major observation is the fact that the buffer level on average is around 40 s, for the conservative switching behavior as well as more aggressive values up to $\alpha = 0.6$. Larger buffer levels reduce the risk of stalling events due to wrong adaptation decisions. From this it follows that an adaptation logic can prefer higher average quality and still keep a comfortable buffer level during playback.

Subsequently, we take a look at the tradeoff between the average quality and the switching frequency as the median of the different starting timestamps for the video CRZbG73SX3s for the medium goodput pattern. Figure 6 shows the average playback quality (left axis) and the switching frequency (right axis) for different values of α.

For $\alpha \leq 0.07$, the number of switches is zero and the average playback quality is 3.0, as also observed in Figure

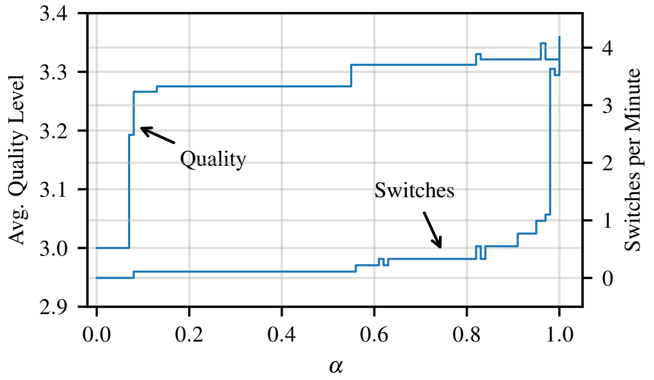

Figure 6. Average playback quality versus switching frequency for video CRZbG73SX3s for the medium goodput pattern as median over the pattern starting timestamps. Starting from $\alpha = 0.1$ a more aggressive switching strategy does not result in increasing average playback quality.

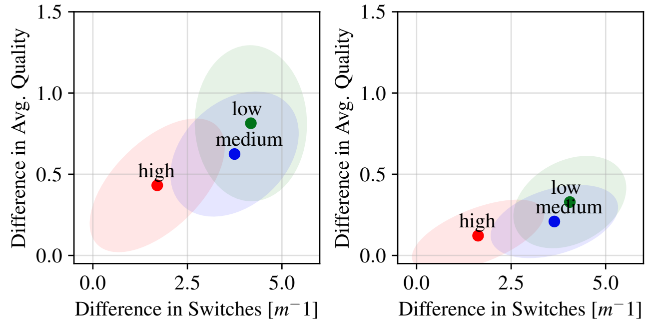

(a) $\alpha = 0.01$ (b) $\alpha = 0.1$

Figure 7. Difference in average quality and switching rate between $\alpha = 1.0$ and $\alpha = \{0.01, 0.1\}$ for the three goodput patterns low, medium and high. The shaded areas denote the 2 dimensional standard deviation of the samples.

4. For values of α between 0.07 and 0.55, the switching frequency increases to 0.1 switches per minute and the quality increases rapidly to 3.28. The difference in average quality compared to Figure 4 is due to the fact that we consider here all starting timestamps of the goodput pattern, while Figure 4 shows only one particular. Starting from $\alpha = 0.6$ to $\alpha = 0.94$, the switching frequency increases up to 0.8 switches per minute, while the average quality stagnates at around 3.31. If α is further increased, the switching rate increases rapidly up to 4.2 switches per minute while the average quality only increases marginally to 3.32. This example is in line with the previous observations that a more aggressive switching frequency does not necessarily benefit the average playback quality. In contrary, the experience of the user is diminished by frequent quality switches while on average the playback quality can not be increased by the frequent switches.

B. Observations for all videos

Next we evaluate the following question by looking at the whole set of videos. *What is the maximum achievable gain in terms of average playback quality when using an aggressive switching strategy compared to a conservative one?* Figure 7 presents the difference in switches rate and difference in average playback quality between $\alpha = 0.01$ and $\alpha = 1.0$ (Fig. 7(a)) and between $\alpha = 0.1$ and $\alpha = 1.0$ (Fig. 7(b)) for the three goodput patterns over all 41 videos. The shaded areas denote the 2 dimensional standard deviation of the samples. The dots represent the mean of the samples.

Multiple conclusions can be drawn from the figure. First, the lowest gain in quality can be observed for the high goodput pattern. This is due to the fact that there are many videos in the set where the high goodput pattern offers enough traffic volume to download always the highest quality level. There are no switches needed for those videos. Second, the mean of the low goodput pattern reaches $4.6\,m^{-1}$ switching rate and a difference of 0.8 quality level for $\alpha = 0.01$. This means that on average in a low goodput scenario you can improve the average quality level by 0.8 by aggressively switching the quality level on average every 13 s. Third, the difference in average quality drops considerable when comparing $\alpha = 0.1$ with $\alpha = 1.0$. For example, the maximum achievable gain for the low goodput pattern drops from 0.8 to about 0.4. Consequently, from the user's perspective there is only a small gain in switching with a higher aggressiveness than $\alpha = 0.1$.

VI. Summary & Discussion

We first summarize the methodology and the findings of the evaluation. Afterwards we discuss the implications of the findings and the application of this paper to the future work in this area.

At first we discuss the methodology and we propose a modification of an existing optimization formulation which allows us to calculate an optimal adaptation path for a given video, a given goodput pattern and a given switching-versus-quality trade-off parameter. This trade-off parameter defines the aggressiveness of the quality switching behavior and is denoted by α. Choosing $\alpha = 0$ results in an adaptation path with zero quality switches and $\alpha = 1$ results in an adaptation path which tries to optimize the average quality at all cost, i.e. with as many quality switches as necessary. Based on an example video we observe that the average buffer level drops fast for values of α between 0.0 and 0.1. For $\alpha > 0.1$, the average buffer level stays close to 42 s and starting from $\alpha = 0.6$ drops linearly to 20 s. Afterwards we take a look at the average playback quality and the switching frequency for the example video. The results show that the switching frequency increases with $\alpha \geq 0.6$ rapidly, while the average quality increases fast and reaches its maximum early at about $\alpha = 0.1$. An evaluation of 41 randomly selected videos from YouTube shows that on average an increase of up to one quality level is possible by increasing the switching frequency by 5 switches per minute. However, the evaluation also shows that this increase is half due to the sharp increase in average quality for $\alpha = 0.01$ to $\alpha = 0.1$.

In general the results show that aggressive switching behavior is not necessary rewarded with a higher average playback quality. From the evaluation also follows that a good starting point for future evaluation of the α parameter is $\alpha = 0.1$.

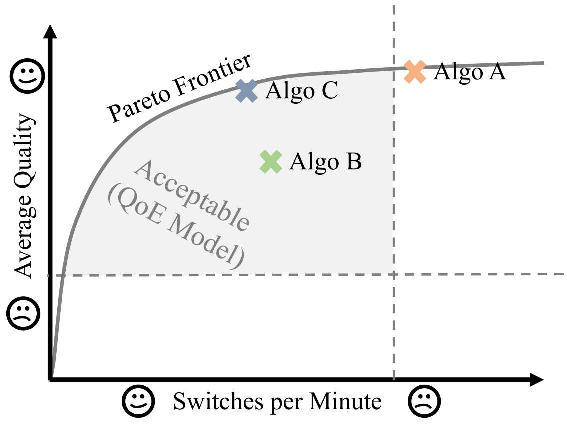

Figure 8. Application of our contribution for future work. The optimization formulation enables us to determine the Pareto frontier for any given HAS adaptation algorithm, a given goodput pattern and a given video in terms of the trade-off between average quality and switching frequency. In combination with a future sophisticated QoE model, this allows for a novel evaluation and classification of adaptation algorithms for HAS video streaming.

Higher, more aggressive, values do not increase the average playback quality much further.

Figure 8 illustrates qualitatively the application of our contribution to the future work in the area of HAS adaptation research. The optimization formulation allows us to determine the Pareto frontier for any given HAS adaptation considering the trade-off between maximizing the average quality and minimizing the number of quality switches during playback. This means that no existing or future HAS adaptation can reach a higher average quality for a given number of switches than the Pareto frontier. Furthermore, we know from the evaluation that the Pareto frontier quickly saturates in terms of average quality. It is ongoing user-experience research to determine a model for the lower bound for the average playback quality and an upper bound for the switching frequency. In combination with such a future sophisticated QoE model, the Pareto frontier allows for a novel evaluation and classification of adaptation algorithms for HAS video streaming.

VII. Conclusion

In this paper, we investigated the trade-off between the average quality and the number of quality switches in a video streaming session. For this purpose we use a quadratic program that includes these two values in its optimization function in order to receive optimal values.

Our results show that high average quality can be achieved with few switches while a very high number of switches is necessary to achieve the highest possible quality. We conclude that it is advisable to rely on conservative strategies that do not switch with a high frequency. However, it is still subject of future work to what degree the number of switches actually impacts the QoE. In future work, we plan to determine the Pareto frontier for the trade-off between the average quality and the number of quality switches for HAS.

Acknowledgements

This work was partly funded by Deutsche Forschungsgemeinschaft (DFG) under grants HO 4770/1-2 (DFG Ökonet: Design and Performance Evaluation of New Mechanisms for the Future Internet – New Paradigms and Economic Aspects).

References

[1] M. Zink, J. Schmitt, and R. Steinmetz, "Layer-encoded video in scalable adaptive streaming," *IEEE Transactions on Multimedia*, vol. 7, no. 1, pp. 75–84, 2005.

[2] S. Egger, B. Gardlo, M. Seufert, and R. Schatz, "The impact of adaptation strategies on perceived quality of http adaptive streaming," in *Proceedings of the 2014 Workshop on Design, Quality and Deployment of Adaptive Video Streaming*. ACM, 2014, pp. 31–36.

[3] T.-Y. Huang, R. Johari, and N. McKeown, "Downton abbey without the hiccups: Buffer-based rate adaptation for http video streaming," in *Proceedings of the 2013 ACM SIGCOMM workshop on Future human-centric multimedia networking*. ACM, 2013, pp. 9–14.

[4] J. Jiang, V. Sekar, and H. Zhang, "Improving fairness, efficiency, and stability in HTTP-based adaptive video streaming with festive," in *Proceedings of the 8th international conference on Emerging networking experiments and technologies*. ACM, 2012, pp. 97–108.

[5] Z. Li, X. Zhu, J. Gahm, R. Pan, H. Hu, A. C. Begen, and D. Oran, "Probe and adapt: Rate adaptation for HTTP video streaming at scale," *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 4, pp. 719–733, 2014.

[6] J. Kua, G. Armitage, and P. Branch, "A survey of rate adaptation techniques for dynamic adaptive streaming over http," *IEEE Communications Surveys & Tutorials*, 2017.

[7] C. Timmerer, M. Maiero, and B. Rainer, "Which adaptation logic? an objective and subjective performance evaluation of http-based adaptive media streaming systems," *arXiv preprint arXiv:1606.00341*, 2016.

[8] J. Añorga, S. Arrizabalaga, B. Sedano, J. Goya, M. Alonso-Arce, and J. Mendizabal, "Analysis of youtubes traffic adaptation to dynamic environments," *Multimedia Tools and Applications*, pp. 1–24, 2017.

[9] M. Seufert, S. Egger, M. Slanina, T. Zinner, T. Hoßfeld, and P. Tran-Gia, "A survey on quality of experience of http adaptive streaming," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 1, pp. 469–492, 2015.

[10] P. Juluri, V. Tamarapalli, and D. Medhi, "Measurement of quality of experience of video-on-demand services: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 401–418, 2016.

[11] L. Yitong, S. Yun, M. Yinian, L. Jing, L. Qi, and Y. Dacheng, "A study on quality of experience for adaptive streaming service," in *Communications Workshops (ICC), 2013 IEEE International Conference on*. IEEE, 2013, pp. 682–686.

[12] T. Hoßfeld, M. Seufert, C. Sieber, and T. Zinner, "Assessing effect sizes of influence factors towards a qoe model for http adaptive streaming," in *Quality of Multimedia Experience (QoMEX), 2014 Sixth International Workshop on*. IEEE, 2014, pp. 111–116.

[13] F. Metzger, E. Liotou, C. Moldovan, and T. Hoßfeld, "Tcp video streaming and mobile networks: Not a love story, but better with context," *Computer Networks*, vol. 109, pp. 246–256, 2016.

[14] E. Liotou, T. Hoßfeld, C. Moldovan, F. Metzger, D. Tsolkas, and N. Passas, "Enriching http adaptive streaming with context awareness: A tunnel case study," in *Communications (ICC), 2016 IEEE International Conference on*. IEEE, 2016, pp. 1–6.

[15] T. Hoßfeld, M. Seufert, C. Sieber, T. Zinner, and P. Tran-Gia, "Identifying qoe optimal adaptation of http adaptive streaming based on subjective studies," *Computer Networks*, vol. 81, pp. 320–332, 2015.

[16] C. Moldovan, C. Sieber, P. Heegaard, W. Kellerer, and T. Hoßfeld, "Youtube can do better: Getting the most out of video adaptation," in *Teletraffic Congress (ITC 28), 2016 28th International*, vol. 3. IEEE, 2016, pp. 7–12.

[17] V. Joseph, S. Borst, and M. I. Reiman, "Optimal rate allocation for adaptive wireless video streaming in networks with user dynamics," in *INFOCOM, 2014 Proceedings IEEE*. IEEE, 2014, pp. 406–414.

[18] "Supplemental material to the paper." [Online]. Available: github.com/csieber/alpha-dataset

[19] C. Sieber, P. E. Heegaard, T. Hoßfeld, and W. Kellerer, "Sacrificing efficiency for quality of experience: YouTube's redundant traffic behavior," in *IFIP Networking 2016 Conference (Networking 2016)*, Vienna, Austria, May 2016.