

A fair share for all: novel adaptation logic for QoE fairness of HTTP adaptive video streaming

Michael Seufert, Nikolas Wehner, Pedro Casas, Florian Wamser

Angaben zur Veröffentlichung / Publication details:

Seufert, Michael, Nikolas Wehner, Pedro Casas, and Florian Wamser. 2018. "A fair share for all: novel adaptation logic for QoE fairness of HTTP adaptive video streaming." In 2018 14th International Conference on Network and Service Management (CNSM), 5-9 November 2018, Rome, Italy, edited by Stefano Salsano, 19-27. Piscataway, NJ: IEEE. <https://ieeexplore.ieee.org/document/8584996/>.



A Fair Share for All: Novel Adaptation Logic for QoE Fairness of HTTP Adaptive Video Streaming

Michael Seufert*, Nikolas Wehner*, Pedro Casas*, Florian Wamser†

*AIT Austrian Institute of Technology GmbH, Vienna, Austria

†University of Würzburg, Würzburg, Germany

michael.seufert.fl@ait.ac.at, nikolas.wehner@ait.ac.at, pedro.casas@ait.ac.at, florian.wamser@informatik.uni-wuerzburg.de

Abstract—This paper presents a novel adaptation logic for HTTP adaptive streaming (HAS), which achieves not only a high Quality of Experience (QoE) but also high QoE fairness among independent and heterogeneous clients. The algorithm forces video clients to adapt the requested quality level based on the current network conditions and their individual bit rate requirements, such that the overall quality levels selected by all currently active streaming clients are fairly distributed, i.e., they do not diverge too much. The design of the algorithm is inspired by the well-known Transmission Control Protocol (TCP) congestion control, and drives heterogeneous clients to independently converge on similar quality levels without the need for communicating with each other and/or with a centralized controller in the network. By defining quality levels with equal visual quality, and preparing video representations accordingly, the quality level fairness is extended to QoE fairness. In this work, the design of the algorithm is described and a simulative performance evaluation is conducted to compare the QoE and QoE fairness of the proposed algorithm with other HAS adaptation logics.

I. INTRODUCTION

Video services on the Internet have evolved from offering pure downloads of video files to progressive downloads and streaming, which both describe the concurrent download and playback of media files. Today, streaming services are the most popular and most demanding applications of the Internet due to the high number of requests, high bit rates of the video content and strict real-time requirements of the video playback. Still, the delivered streaming service has to meet the expectations of the end users. To understand and eventually improve Internet services like video streaming, application providers and Internet service providers use the concept of Quality of Experience (QoE) to quantify the subjective experience and satisfaction of their customers with the network and the service. For video streaming, the most severe QoE degradations were the waiting time until the start of the playback (initial delay) and interruptions of the playback (stalling) [1]–[3].

The currently prevailing streaming technology – HTTP Adaptive Streaming (HAS) – allows to mitigate these QoE degradations by offering the possibility to adapt the video bit rate to the network conditions. The goal is to ensure a smooth streaming when end users face throughput fluctuations, e.g., in mobile networks. HAS utilizes standard web protocols (mostly

HTTP over TCP, recently also over QUIC) to promote a simple service implementation and high availability. It is implemented in many commercial solutions and was standardized as MPEG Dynamic Adaptive Streaming over HTTP (DASH) [4].

To enable adaptation of the video streaming to the current network conditions, the HAS server stores the video content encoded in different representations, i.e., in different bit rates. The representations are split into segments (also referred to as chunks) and corresponding segments of different representations contain the same frames, such that the bit rate can be seamlessly switched at each segment boundary. Typically, all segments contain a fixed amount (i.e., 1 to 15 s) of video playback time and can either be extracted at runtime from the single representation file or are stored as separate files on the server. To change the video bit rate, the transmitted video has to be altered, e.g., in terms of resolution, frame rate, or compression, which changes the visual quality of the segment, thereby, introducing an additional impact on QoE [1], [5]–[7].

The adaptation logic at the client is an algorithm that selects which segments to download next from the list of available segments. These decisions usually take into account segment characteristics (e.g., bit rate), current playout statistics (e.g., buffer fill level), and current network conditions (e.g., bandwidth measurements or estimations). By adapting the video bit rate appropriately, the algorithm aims to maximize the QoE for the given network conditions by minimizing initial delay, avoiding stalling, and playing out the video in a high visual quality. Consequently, it is the adaptation logic, which predominantly influences the network demands of the video streaming and the resulting QoE of the end user.

Network providers want the resulting QoE of the end users in their networks to be high, but also similar, which can be quantified in terms of QoE fairness [8]. When multiple HAS flows are in a network and share a bottleneck link, the link capacity will be equally distributed to all flows by the fairness of the TCP protocol, and thus, also the bit rates of the videos will be similar. However, as the requirements for each streaming flow might be different, e.g., in terms of content complexity or device characteristics, the resulting QoE of the end users might not be fair at all. For example, smartphones that stream a video in low resolution only need a small bit rate to achieve a high visual quality (little compression) and high QoE. However, streaming the same video on a TV in large resolution with the same bit rate will result in a low

visual quality (much compression) due to the higher resolution, and thus a lower QoE. This shows that, while the bit rate distribution in the network is fair, the QoE distribution is not fair at all. Existing HAS adaptation logics were not able to cope with the problem of QoE fairness for competing HAS flows, such that collaborative and centralized solutions in the network were needed [9]–[12].

This paper presents a novel HAS adaptation logic, which reaches high QoE as well as QoE fairness on shared links with multiple HAS flows without requiring communication between the streaming applications and/or a centralized controller. It is inspired by the well-known Transmission Control Protocol (TCP) congestion control, and drives heterogeneous clients to independently converge on similar quality levels, instead of similar bit rates. When equal quality levels correspond to equal visual quality, which can be easily realized during the content preparation by the video provider, the QoE fairness in the network will be high, i.e., all HAS flows will achieve a similar QoE by using the proposed adaptation logic.

The remainder of the paper is structured as follows. Section II describes related works on adaptation logics for HAS and QoE fairness. Moreover, it recaps the basic principles of TCP, especially focusing on its congestion control to achieve fairness. Section III presents the design considerations for the novel HAS adaptation logic and shows the implementation of the algorithm in detail. A performance evaluation of the proposed algorithm is conducted in Section IV, in which it is compared to existing HAS adaptation logics in terms of QoE and QoE fairness, and Section V concludes.

II. RELATED WORK

The Quality of Experience (QoE) of HTTP adaptive streaming (HAS) is a widely investigated research field. The most important research results were summarized in [1], while more recent publications confirmed the findings that initial delay, stalling, and quality adaptation are the most dominant QoE factors. While the initial delay has only a small impact on the QoE, stalling, i.e., the interruption of the video playback due to buffer depletion, is considered the worst form of QoE degradation [2], [3]. Moreover, the played out video quality and the time on each quality layer strongly impact the QoE [5]–[7]. These QoE factors will be considered during the performance evaluation of the proposed algorithm.

As subjective QoE studies, which are typically used to assess the QoE, are expensive and time-consuming, QoE models are needed, which can predict the QoE based on objective criteria, e.g., the QoE factors. Recently, such an objective QoE model was standardized as P.1203 [13], whose mode 3 (i.e., complete information is available) will also be used in this paper to obtain QoE scores for HAS sessions. The output of the P.1203 metric is a predicted mean opinion score (MOS) on a 5-point absolute category rating (ACR) scale, ranging from 1 (bad) to 5 (excellent) [14]. The predicted score is shaped by information about the audio and video encoding, as well as by application-layer parameters like the number of stallings, the length of the stalling events, the initial

delay, and the number of quality switches. More details on the implementation of the standard can be found in [13].

QoE factors are not only directly influenced by the network conditions, but also by the adaptation logic of the HAS client. It can be considered a mechanism for application management, which aims to maximize the user’s QoE by requesting appropriate quality levels in order to minimize initial delay, avoid stalling, and download the video with a high visual quality. Network providers do not only want to maximize the QoE of each end user, but also want to achieve a high QoE fairness in their networks. Therefore, they might apply network management or collaborative application and network management. However, as these solutions are not practical and require business agreements between service providers and network providers, the goal of this work is to reach QoE fairness of independent and heterogeneous HAS clients by only application management. Therefore, a new adaptation logic for HAS clients is proposed, which is inspired by TCP fairness and aims to both maximize QoE and QoE fairness.

a) HAS Adaptation Logics: Literature proposes several HAS adaptation logics that focus on different aspects. Some of them are considered in this paper, and will be shortly described in the following. BufferBased [15] solely utilizes the buffer level to decide the next quality level. By dividing the buffer level into several segments and assigning different actions, e.g., increasing/decreasing the quality level, the adaptation logic maps the current buffer level to the segments, and performs the corresponding action. In contrast, ELASTIC [16] applies linear feedback control theory in order to avoid the ON-OFF traffic pattern that occurs when clients estimate their bandwidth shares falsely due to other temporarily inactive clients. KLUDDCP [17] selects the quality level based on the disparity between the current buffer level and the desired buffer level, a bandwidth estimate, and the corresponding bit rates. Finally, TRDA [18] is designed for stalling prevention and reduction of the frequency of quality adaptations.

[19] uses a simulation framework based on the Wi-Fi model of NS-3 to compare several HAS adaptation logics, namely, TRDA/Tobasco, Conventional, PANDA [20], and FESTIVE [21]. They found that TRDA showed the best results with respect to the stalling ratio, the number of quality changes, and the average buffer level, but suffered from the lowest mean video quality. The best video quality was obtained by FESTIVE, which in turn suffered from frequent quality switches and, for too many clients, stalling. [22] compares ten different adaptation logics and considers QoE metrics, as well as network utilization, stability of the adaptation, and the mean bit rates requested by the algorithms. The results revealed no dominating adaptation logic, but even simple adaptation logics performed reasonably well.

b) QoE Fairness: QoE fairness [8] refers to a fair distribution of the QoE scores of all end users in a network. This means that the QoE scores should be very similar and not diverge too much. [8] found that the prevailing fairness measure, Jain’s fairness index [23], was not particularly well suited to quantify QoE fairness. The reason is that Jain’s

fairness index requires ratio scales with a defined zero point, but QoE is rather measured on interval scales, e.g., the 5-point ACR scale, which is typically used to obtain MOS values.

[8] defined a new measure for QoE fairness F with additional scale and metric independence for computing the fairness of QoE scores. Thereby, the fairness measure F is computed as a linear transformation of the standard deviation of the QoE scores to the interval [0;1]. This is achieved by setting the standard deviation in relation to the maximum possible standard deviation, which is the difference between the highest possible QoE score H and the lowest possible QoE score L : $F = 1 - \frac{\sigma}{\sigma_{max}} = 1 - \frac{2\sigma}{H-L}$

Consequently, $F = 1$ indicates that $\sigma = 0$ and all QoE scores are the same, i.e., a perfect QoE fairness. In contrast, $F = 0$ represents the most unfair situation, in which 50% of the users have a QoE score of H and 50% have a QoE score of L . In this paper, F is used to compute the fairness of the selected quality levels and the resulting QoE score.

c) QoE Fairness with Additional Network Management:

As current HAS adaptation logics of independent and heterogeneous clients cannot reach QoE fairness in the network, the research community tackled this issue with the help of network management. Thereby, the network conditions are influenced during the video streaming depending on the current QoE, e.g., by prioritization or dedicated bandwidth allocation for video flows [24], [25]. However, this requires the estimation of video requirements and the monitoring of QoE in the network, which has become a difficult and cumbersome task since video traffic is encrypted. Another approach is combined network and application management based on the collaboration and communication between the HAS clients and the network, e.g., by using a centralized controller in the network. In the following, several works are presented.

[26] uses SDN to dynamically allocate the network resources for each client based on its expected QoE. [27] modifies the HAS index files, i.e., the list of available segments, in the network. [28] proposes bandwidth reservation for HAS flows and signals clients which quality levels to request. [9] tries to achieve QoE fairness for multiple clients by using information about network conditions. [10] presents a collaborative traffic management system, which considers QoE fairness in case of encrypted HAS. Recently, SAND standardization efforts were started to exchange information measured at servers and network elements to improve the delivery, and to send quality-related information to clients to improve the reception [11], [12].

d) *TCP Fairness*: The Transmission Control Protocol (TCP) [29] is one of the most widely used protocols on the Internet. TCP is connection-oriented and ensures a guaranteed transmission of undisturbed data. It provides a principle of fair bandwidth shares among multiple clients over a shared network link. This is achieved by utilizing flow control and congestion control. As TCP flow control is rather focused on avoiding to overload the receiver, it was irrelevant for the design of the adaptation logic, and thus, is not discussed here.

The basic TCP congestion control [30] is composed of an

exponential growth of the congestion window, the so called slow start phase, followed by a linear growth of the congestion window, the congestion avoidance phase. The congestion window indicates the number of unacknowledged TCP segments, which the sender can send within one transmission round. The duration of a transmission round is characterized by the round trip time. During the slow start phase, the congestion window is increased by one for each received acknowledgment, resulting in an exponential growth of the window, as the number of segments doubles in each round.

A threshold specifies the end of the slow start phase and hence the beginning of the congestion avoidance phase. In the congestion avoidance phase, the congestion window is increased by one per transmission round. When a segment is lost, congestion occurs. As a consequence, standard TCP Tahoe decreases the congestion window size to 1 and sets the slow start threshold to half of the previously active congestion window. This is followed by a new slow start phase. Several variations of standard TCP exist, which modify this basic procedure. For example, in TCP Reno the congestion window size is not reduced to 1 on loss, but only to half of the previous window size and a congestion avoidance phase directly follows instead of a slow start phase (fast recovery).

As all clients can detect congestion rather simultaneously by observing their own loss, and react similarly by reducing the own sending rate, this can be considered an altruistic behavior of clients. It reaches that the bandwidth of the shared link can be fairly shared among all clients. The proposed HAS adaptation logic was inspired by the TCP congestion control, and uses similar design considerations to achieve a fair quality level for all HAS clients. Note that practical problems might arise with HAS flows in a shared TCP network due to the on-/off-phases of the download [1], [31]. However, these issues are not considered in this work, as it only focuses on investigating the potential of a TCP-inspired adaptation logic.

III. TCP-INSPIRED ADAPTATION LOGIC

This section introduces the novel HAS adaptation logic, which reaches QoE fairness on shared links with multiple independent and heterogeneous HAS clients.

A. Design Considerations

A prerequisite for achieving similar QoE on heterogeneous HAS clients is that each client can request the same quality levels, which is not the case when they can only request different bit rate levels. Typically, different videos with the same bit rate have a different visual quality, e.g., due to different content complexity, resolution, or compression. Therefore, quality levels are introduced. Without loss of generality, the quality levels are ordered from 1 (lowest quality) to l_{max} (highest quality) and the visual quality contained in different segments of the same quality level should be the same. In practice, this means that instead of encoding the different representations of a HAS video file with certain (target) bit rate levels, the target for encoding should be a metric for the visual quality (e.g., SSIM or VQM), such that different

segments of the same quality level have the same visual quality. Video providers also have to ensure that the quality levels are consistent for different contents or resolutions.

The main concept for reaching QoE fairness among independent and heterogeneous HAS clients is based on the idea that all clients need to obtain the same information from observing the network. This information can then be used by each client together with information about its individual characteristics to maximize its own QoE in a synchronous and consistent, yet altruistic way. In this case, the term ‘‘altruistic’’ means that each client is aware that the network is shared and does not egoistically increase its own quality level in situations when other clients can not. In the following, the four main design considerations are described that are used to obtain ‘‘altruistic’’ behavior:

a) Limitation of buffer filling rate: One problem of competing heterogeneous clients is that clients with small bit rates can download segments and fill the buffer more quickly than clients with large bit rates. Consequently, with most adaptation logics, they can much earlier attempt to increase the quality level, which can result in an unfair QoE distribution. In order to keep the clients and their buffer more synchronized, a limitation of the buffer filling rate is introduced. After a desired buffer level b_d is reached, the buffer will be kept at a constant level in terms of buffered segments. Therefore, after a segment is downloaded, the clients have to go idle for the contained playtime of a segment τ minus the download time of the segment T , i.e., $t_i = \tau - T$. Thus, a HAS client with sufficiently filled buffer will request segments regularly in intervals of τ , which results in a buffer filling rate of 1 that is equal to the playback rate (unit rate constraint). Thereby, the buffer is effectively kept constant at the desired buffer level b_d , which also avoids excessive amounts of unnecessarily downloaded data in case of video abortion.

If the buffer is below the desired level b_d , the buffer has to be increased. Therefore, the buffer filling rate can be increased to 2. This is achieved by allowing HAS clients to go idle only for $t_i = \frac{\tau}{2} - T$ after a segment download, which causes them to request segments in intervals of $\frac{\tau}{2}$. As the buffer filling rate is double the playout rate, the buffer will increase linearly with rate 1. If the video is initially loading or stalling, the buffer has to be increased as fast as possible to minimize the initial delay or stalling time. Therefore, in this case, the HAS clients do not need to go idle after a segment download, but can request segments back to back until the playback can be started.

b) Capacity estimation and upper bounds on quality level: After the timing of segment requests is more synchronized among all clients, the requested quality levels have to be synchronized next. Therefore, a common information about the network is needed, which can be observed by each client individually, and can be used to come to an agreement on the quality level. Such information is the maximum capacity \hat{C}_{max} , which is observed by the client. For each of k active HAS clients, its observed \hat{C}_{max} should be similar and approximate $C_{max} = \frac{C^*}{n}$, which is the fair bandwidth share of the link capacity C^* as a result of TCP fairness, when n flows

(k video flows and $n - k$ other flows) share the link. After each segment download, a capacity estimate $\hat{C} = \frac{s}{T}$ can be computed by dividing the size s of the requested segment by the download time T , and the observed maximum capacity \hat{C}_{max} can be updated. All adaptation decisions will be based on the observed maximum capacity. Therefore, a fictitious reference video v' is considered with a maximum needed bit rate r'_{max} equal to the maximum capacity, i.e., $r'_{max} = \hat{C}_{max}$. This reference video would just be able to maintain its buffer when its throughput is equal to the maximum capacity, i.e., the download time of a segment is equal to the playtime contained in a segment. To account for bit rate variations, the maximum needed bandwidth is computed as the mean bit rate of all segments of the highest quality level \bar{r}_{max} plus α -times the corresponding standard deviation $\sigma_{r_{max}}$, i.e., $r_{max} = \bar{r}_{max} + \alpha \cdot \sigma_{r_{max}}$, $\alpha \geq 0$.

If the throughput of a video decreases to a share $0 < q \leq 1$ of the maximum capacity, i.e., $\hat{C} = q \cdot \hat{C}_{max}$, the reference video would have to decrease its quality level to $q \cdot l_{max} := l_u$, and thus, to be fair, also all other videos should reduce their quality levels to l_u . The computation of l_u utilizes a parameter $\beta \leq 1$ to relax the normalization with the individually observed \hat{C}_{max} , i.e., $l_u = \lceil \frac{\hat{C}}{\beta \cdot \hat{C}_{max}} \cdot l_{max} \rceil$. This means that the maximum quality level can still be requested when $\hat{C} \geq \beta \cdot \hat{C}_{max}$. Thus, by obtaining the same (or similar) \hat{C}_{max} and \hat{C} , competing clients can individually find the same (or a similar) upper bound l_u for the quality level of requested segments based on the current network conditions.

c) Slow start and congestion avoidance phase: Similar to TCP, the quality level of HAS clients should be increased fast until the currently supported quality level is reached. Therefore, an exponential growth of the current quality level l_c was implemented similar to the TCP slow start phase. A slow start phase is entered always after the playback of the video started, i.e., after initial delay or stalling. It allows to double the requested quality level after each downloaded segment, i.e., the new quality level $l_n = 2 \cdot l_c$. After the current segment has exceeded half of the currently allowed maximum level l_u , and thus, another doubling is not possible, the congestion avoidance phase follows.

The congestion avoidance phase approaches and aims to reach the maximum supported quality level. In this phase, the quality level can be increased by 1 after each γ segments. Therefore, each client remembers its last quality change t_{l_c} , and if more time than $\gamma \cdot \tau$ has passed, the new quality level $l_n = l_c + 1$ can be requested. Note that the upper bound l_u overrides increases of the quality levels that would have been triggered by the slow start or congestion avoidance phase.

d) Congestion detection and compensation of bit rate requirements: A ‘‘congestion’’ with the current quality level l_c occurs when the download time T of a segment is larger than the contained playtime τ , and thus, the buffer decreases. In this case, l_c should be decreased in order to avoid further depletion of the buffer, which could eventually cause stalling. To consistently act among all HAS clients, they will compute the same new quality level $l_n = \delta \cdot l_c$, being a fixed share

$\delta < 1$ of the current quality level on which the congestion was detected. Moreover, a hard threshold for the buffer level b_l is introduced, which also triggers a congestion. If the buffer drops below the low threshold b_l , the buffer should increase, and thus, the quality level has to be decreased.

As videos with larger bit rates are more likely to face congestion, the segment download times T have to be compensated according to the bit rate requirements. Therefore, again the fictitious reference video v' is considered. Remember, to maintain its buffer, it should be allowed to fully use the download time $T' = \tau$ to download a segment of the highest quality level under the maximum network conditions \hat{C}_{max} . To have a comparable and fair congestion threshold, a video with a smaller maximum needed bit rate r_{max} should have to download the video with a proportionally shorter download time. Thus, the bit rate compensation can be done via the compensated download time T_c , which can be computed from the actual download time T as $T_c = T \cdot \frac{\hat{C}_{max}}{r_{max}}$ and compared to the segment playtime τ to detect congestion. Again $r_{max} = \bar{r}_{max} + \alpha \cdot \sigma_{r_{max}}$ is the individually needed maximum bit rate of each HAS client.

As congestion is always an indicator that the requested quality level cannot be maintained, the assumptions on the network conditions have to be reevaluated. To account for new clients, which reduce the fair bandwidth share $C_{max} = \frac{C^*}{n}$, the assumed maximum capacity \hat{C}_{max} is halved. This is a conservative reaction, as for each additional client (n increases by 1), the new C_{max} is at least half of the previous C_{max} . Still, \hat{C}_{max} can be updated with subsequent capacity estimations.

B. Algorithm Description

The algorithm is called TCP-inspired Adaptaion Logic (TCPAL) and is shown in Algorithm 1. TCPAL has ten predefined parameters on which all HAS clients have to agree, namely, the segment playtime τ , maximum quality level l_{max} , the initial delay threshold b_i , the stalling threshold b_s , the low buffer threshold b_l , the desired buffer level b_d , the considered bit rate variation α , the normalization relaxation β , the linear growth regulation γ , and the congestion decrease factor δ . Additionally, each HAS client has to maintain and accordingly update the maximum needed bit rate r_{max} , the current quality level l_c , the current buffer level B , the time since the last quality change t_{lc} , the currently observed maximum capacity \hat{C}_{max} , and the Boolean variables `isInitialDelay`, `isStalling`, and `slowStart`. These client variables are initialized as shown in ‘‘Client Vars’’ and they are updated either by the algorithm when indicated by ‘‘:=’’ (in contrast to ‘‘=’’, which represents an allocation to a temporal variable), or externally (reduction of B during playback, $B == 0$ (playback is stalling) \Rightarrow `isStalling := True`), or never (r_{max}).

After a video request, the HAS clients will immediately download the first segment in the lowest quality level, i.e., $l_c = 1, t_i = 0$, and subsequently call TCPAL whenever a segment is completely downloaded. Thereby, the only changing inputs are the size of the last downloaded segment s and the download time of the last segment T . TCPAL will output the

quality level of the next segment l_n and the idle time t_i , i.e., the time until the next segment with quality level l_n can be requested. The presented design considerations can be found in Algorithm 1 as follows: The limitation of the buffer filling rate is implemented in Lines 2-15. The capacity estimation and upper bound to the available quality levels can be seen in Lines 16-18. The congestion detection in Lines 19-20 is followed by the corresponding consequences of congestion in Lines 21-23. The slow start phase corresponds to Lines 25-28, and the congestion avoidance phase to Lines 29-35.

```

Parameters      :  $\tau, l_{max}, b_i, b_s, b_l, b_d, \alpha, \beta, \gamma, \delta$ 
Client Vars    :  $r_{max} := \bar{r}_{max} + \alpha \cdot \sigma_{r_{max}}, l_c := 1, B := 0,$ 
                    $C_{max} := 0, t_{lc} := t_{NOW}, isInitialDelay := True,$ 
                    $isStalling := True, slowStart := True$ 

Input:  $s, T$ 
Output:  $l_n, t_i$ 

1  $B := B + \tau;$ 
2 if isStalling == True then
3    $t_i := 0;$ 
4   slowStart := True;
5   if (isInitialDelay == True and B  $\geq$  bi) or (isInitialDelay == False and
6      $B \geq b_s$ ) then
7     playVideo();
8     isInitialDelay := False;
9     isStalling := False;
10  else
11    if  $B \geq b_d$  then
12       $t_i = \max(\tau - T, 0);$ 
13    else
14       $t_i = \max(\frac{\tau}{2} - T, 0);$ 
15    end
16   $\hat{C} := \frac{\hat{C}}{2};$ 
17   $\hat{C}_{max} := \max(\hat{C}, \hat{C}_{max});$ 
18   $l_u = \max(\min(\lceil \frac{\hat{C}}{\beta \cdot C_{max}} \cdot l_{max} \rceil, l_{max}), 1);$ 
19   $T_c = T \cdot \frac{\hat{C}_{max}}{r_{max}};$ 
20  if  $T > \tau$  or  $T_c > \tau$  or  $B - t_i < b_l$  then
21     $l_n = \max(\min(\delta \cdot l_c, l_u), 1);$ 
22     $\hat{C}_{max} := \frac{\hat{C}_{max}}{2};$ 
23    slowStart := False;
24  else
25    if slowStart then
26       $l_n = \min(2 \cdot l_c, l_u);$ 
27      if  $l_n > \frac{l_u}{2}$  then
28        slowStart := False;
29      else
30        if  $t_{NOW} - t_{lc} > \gamma \cdot \tau$  then
31           $l_n = \min(l_c + 1, l_u);$ 
32        else
33           $l_n = \min(l_c, l_u);$ 
34        end
35      end
36    end
37  if  $l_n \neq l_c$  then
38     $t_{lc} := t_{NOW};$ 
39     $l_c := l_n;$ 
40  return  $l_n, t_i$ 

```

Algorithm 1: TCPAL

IV. PERFORMANCE EVALUATION

This section describes the simulation framework, on which the performance evaluation of the proposed adaptation logic is based, and the corresponding results.

A. Simulation Framework

In [25], a Java discrete-event simulation for video streaming and web page downloads over a bottleneck link with fixed capacity was presented. In this work, the simulation was




Fig. 1: Target SSIM values for quality metrics and corresponding bit rates for different resolutions.

extended in order to provide HAS capabilities and to integrate different HAS adaptation logics.

The simulation allocates the available link capacity as follows. Each of the active application flows receives an equal share of the capacity, i.e., the normal TCP-fair share of bandwidth. Whenever a client enters or leaves the system, or the download of a client becomes idle or resumes, the available capacity is redistributed, such that every downloading client is allocated the updated fair bandwidth share. The simulation keeps track of the download and playback of all videos, which can be used to compute several key performance indicators, e.g., the initial delay, stalling, and played out video qualities. Each simulation run is aborted after 550 s and all statistics are stored in log files, which are later used for the evaluations.

For a more realistic simulation, real video chunks were integrated into the simulation. Therefore, the video “Big Buck Bunny” [32] was split into segments with a playback length of four seconds and each segment was encoded with constant bit rate encoding with FFmpeg. In order to obtain quality levels with an equal visual quality, the well-known full-reference metric SSIM (structural similarity) [33] was utilized. 32 target SSIM values were selected as quality levels, ranging from 0.900 (lowest quality level) to 0.995 (highest quality level). Since there are currently no tools that are capable of encoding videos based on a passed target SSIM, the encoding process was performed manually using the method of trial and error, i.e., the encoding bit rate of the segment was varied until the calculated mean SSIM matched the target SSIM. To account for different requirements of HAS clients, e.g., in terms of display characteristics, the video was encoded in three different resolutions (360p, 720p, 1080p), and each client was assigned one of the resolutions at the beginning of a simulation run.

The extracted video representations are depicted in Figure 1, where the y-axis represents the target SSIM values, the x-axis describes the required encoding bit rate in order to obtain the corresponding SSIM value, and each resolution is colored differently. The figure shows that the HD video content requires a much higher bit rate for a similar SSIM compared to the other video sequences.

B. Results


In the following, the performance of the proposed TCP-inspired adaptation logic (TCPAL) is investigated in detail. Therefore, the parameters were set to $\tau = 4$ s, $l_{max} = 32$, $b_i = 16$ s, $b_s = 8$ s, $b_l = 8$ s, $b_d = 16$ s, $\alpha = 1$, $\beta = 0.9$, $\gamma = 2$, and $\delta = 0.75$.

Figure 2 studies the behavior of HAS clients using TCPAL qualitatively in exemplary simulation runs with link capacity 7 Mbps. In Figure 2a, there is only one HAS client on the link, streaming a resolution of 720p. The x-axis depicts the time of the simulation run, while the y-axis shows the requested quality level. In the beginning, the slow start phase can be seen, in which the quality level increases exponentially. The congestion avoidance phase follows, in which the quality level is increased linearly until congestion is detected and the quality level drops to 75% (δ) of the previous level. It can be seen that TCPAL repeatedly approaches the maximum supported quality level, which is 30 given the link capacity of 7 Mbps (cf. Figure 1), and then drops due to congestion.

Figure 2b shows a situation, in which six HAS clients start simultaneously to request the 720p video. It can be seen that all clients act quite synchronous during both slow start and congestion avoidance phase. It can be seen that the maximum supported quality levels are much smaller, because the link capacity is equally shared among the six HAS clients. Figure 2c shows the same situation, but now the six clients request different video resolutions. Two clients request 360p, two clients 720p, and two clients 1080p. This leads to heterogeneous bit rate requirements in the network. It can be seen that the clients are at different quality levels at the end of the slow start phase. However, during the congestion avoidance phase, the clients bring their quality levels to a similar range.


Finally, Figure 2d shows the most challenging scenario, in which six diverse HAS clients (2x360p, 2x720p, 2x1080p) start to request their videos at a uniformly random time in the first 100 s. Thus, the HAS clients are out of sync in this scenario, and the available bandwidth of each client will change when a new HAS client starts to request his video. Nevertheless, it can be seen that the quality levels requested by TCPAL are still quite similar. Later, this scenario will be further investigated to compare the performance of TCPAL to other adaptation logics in terms of QoE and QoE fairness.

To understand the shortcomings of other HAS adaptation logics, Figure 3 shows the behavior of other HAS adaptation logics in exemplary simulation runs of the same scenario of Figure 2d. BufferBased (Figure 3a) increases the quality level quite synchronous for all HAS clients but very slowly. Also TRDA (Figure 3d) shows a linear increase, but as it increases one quality level per segment, high quality levels can be reached much faster. In contrast, ELASTIC (Figure 3b) and KLUDCP (Figure 3c) almost immediately jump to very high quality levels at the start of the video. While the behavior of each adaptation logic is different (e.g., high oscillations of the quality level for KLUDCP, but inert adaptation for ELASTIC or TRDA), all adaptation logics converge to different qual-




(a) Single HAS client with 720p.

(b) Six HAS clients with 720p and simultaneous start.



(c) Six HAS (2x360p, 2x720p, 2x1080p) clients with simultaneous start.



(d) Six diverse HAS (2x360p, 2x720p, 2x1080p) clients with randomized start.

Fig. 2: Performance of TCPAL in different scenarios with link capacity 7 Mbps.

ity levels for different resolutions. This means, they cannot achieve fair adaptation for heterogeneous clients with different bit rate requirements. The quality level, and thus, the QoE of clients with high video bit rates (e.g., 1080p) is much lower than for clients with low bit rate requirements (e.g., 360p).

In the following, the quantitative performance evaluation will be presented. Therefore, the reference scenario of Figure 2d was used, i.e., six clients (2x360p, 2x720p, 2x1080p) can randomly (uniformly distributed) start a video in the first 100s of a simulation run. Three different link capacities are used, namely, the already presented capacity of 7 Mbps, for which the six clients have to find appropriate intermediate quality levels, a low link capacity of 2.4 Mbps, and a high link capacity of 70 Mbps. In case the six clients share a link capacity of 2.4 Mbps, the clients can only support low quality levels and need to avoid stalling. In case of 70 Mbps, all six clients should be able to reach very high quality levels. Note that fluctuations of the link capacity are not investigated. However, the increasing number of HAS clients and their oscillating download and idle phases change the fair bandwidth share of all active HAS client, which can be considered a kind of fluctuation. 20 simulation runs were conducted for each adaptation logic and link capacity. Each simulation run is aborted after 550s, and the performance evaluation (except for initial delay) considers only the steady state phase from 150-550s (dashed lines in Figures 2d and 3).

Figure 4 shows the QoE results for the different link capacities on the x-axes. Each bar represents a different adaptation logic, see the legend of Figure 5 for a mapping

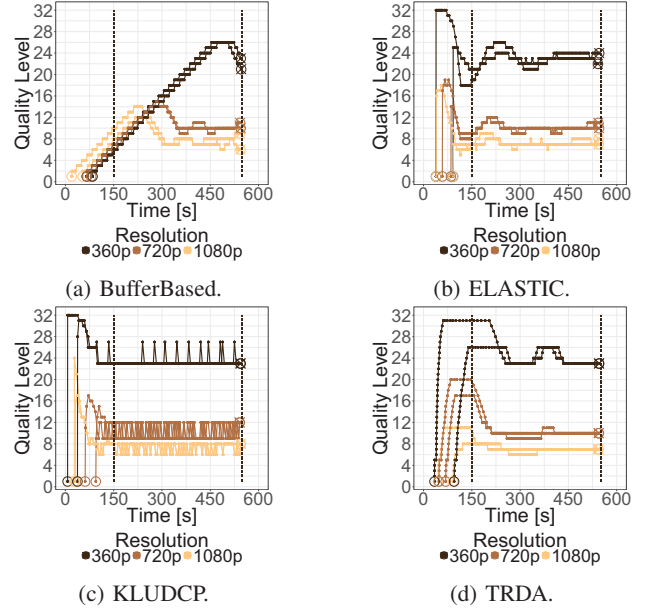


Fig. 3: Performance of other HAS adaptation logics for six HAS (2x360p, 2x720p, 2x1080p) clients with randomized start and link capacity 7 Mbps (cf. Figure 2d).

of colors and algorithms. The y-axes indicate the mean of the considered performance metric, i.e., initial delay (Figure 4a), total stalling time (Figure 4b), time-weighted average quality level (TWAQL, Figure 4c), and the predicted MOS from P.1203 (Figure 4d), and the corresponding 95% confidence intervals are shown on top.

Figure 4a shows the results for initial delay. It can be seen that TCPAL (yellow) reaches low initial delays comparable to TRDA (orange) for all link capacities. Only BufferBased (black) starts the playback of the streamed video faster. The results for stalling are presented Figure 4b. TCPAL cannot completely avoid stalling for the lowest link capacity, however, the mean total stalling time is short. For higher link capacities stalling is not an issue for TCPAL.

Figure 4c displays the time-weighted average quality level, i.e., each quality level was weighted according to the time for which it was played out. TCPAL on average requests slightly lower quality levels than other adaptation logics for higher link capacity. This comes by design as TCPAL shows an altruistic behavior and restricts itself from requesting too high quality levels. Additionally, the congestion avoidance phase, in which the maximum supported quality level is approached linearly and the quality level is reduced on congestion, leads to an average quality level slightly below the maximum supported quality level. This propagates to the MOS prediction based on P.1203 in Figure 4d. It can be seen that ELASTIC (dark brown) and KLUDCP (light brown) can achieve the highest MOS for all link capacities. Nevertheless, TCPAL reaches a comparably high MOS level for 2.4 Mbps and 7 Mbps, and only a slightly lower MOS for 70 Mbps.

Finally, Figure 5 shows the corresponding fairness of the

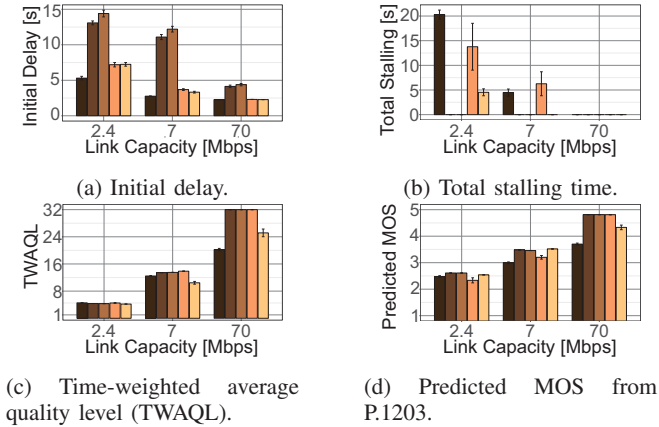


Fig. 4: Results on QoE factors and QoE score.




Fig. 5: Results on QoE fairness.

QoE metrics. While Jain’s fairness index [23] was used for initial delay and stalling (unlimited metrics on ratio scale), the F metric [8] was used for time-weighted average quality level and the predicted MOS scores. The desirable fairness score is 1, which means a perfect fairness of the underlying metric for all six HAS clients.

In case of the low link capacity of 2.4 Mbps, clients with high resolutions struggle with stalling and very low quality levels, while clients with low bit rate requirements can maintain a smooth streaming on low quality levels. Thus, the resulting QoE is very different, which can be seen in the low fairness scores. Also for TCPAL, only the congestion detection via the actual download time works, and thus, clients with low bit rate requirements manage to download segments with a higher quality level as can be seen in Figure 6a. As this is the same problem that other HAS adaptation logics face, the fairness scores of TCPAL are on the same level in this scenario.

For 7 Mbps, stalling should be avoided by all adaptation logics. Still, the fairness of the other adaptation logics is not very high as the requested quality levels highly diverge for different bit rate requirements (cf. Figure 3). In this scenario, TCPAL manages in all simulation runs to reach a significantly higher fairness in terms of time-weighted average quality level (TWAQL) and predicted MOS (P.1203). Together with




Fig. 6: Performance of TCPAL for six HAS (2x360p, 2x720p, 2x1080p) clients with randomized start and smaller or larger link capacity (cf. Figure 2d).

the comparably high overall QoE, this shows that the design considerations of TCPAL were useful and allow a good and fair streaming of heterogeneous and independent HAS clients.

For the high link capacity of 70 Mbps, initial delays are short, no stalling occurs and a constant high quality level can be maintained by the other adaptation logics. By design, the quality level of TCPAL oscillates in the congestion avoidance phase due to the strict definition of congestion, which can be seen in Figure 6b. This should not negatively affect the QoE fairness, however, it can be seen that the requested quality levels additionally diverge. The reason for this problem is that the download times of segments are very small, and thus, the observed maximum link capacity might be different for different HAS clients depending on how many other clients are downloading or idling. This leads to self-restrictions of only some clients based on congestion avoidance. Thus, the requested quality levels are different, which results in reduced fairness for time-weighted average quality level (TWAQL), and also for predicted MOS (P.1203). To improve the performance of TCPAL in this case, the robustness of the estimation of the maximum link capacity would have to be improved.

V. CONCLUSION

This paper presented a TCP-inspired adaptation logic (TCPAL) for QoE fairness among independent and heterogeneous clients of HTTP adaptive video streaming. The design considerations of the adaptation logic are based on the TCP congestion control and drive HAS clients to converge on similar quality levels instead of similar bit rates. The performance evaluation showed that TCPAL performed on par with other HAS adaptation logics in terms of QoE and QoE fairness for low link capacities, but could significantly improve the QoE fairness for increased link capacity. Still, TCPAL has to be improved to perform better for very high capacities, which will be tackled in future works, but the current results are already very promising. Additionally, in future works, the impact of the parameters will be studied, and TCPAL will be compared to more HAS adaptation logics, and evaluated in more scenarios.

REFERENCES

- [1] M. Seufert, S. Egger, M. Slanina, T. Zinner, T. Höbfeld, 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.
- [2] D. Ghadiyaram, J. Pan, and A. C. Bovik, "A Time-varying Subjective Quality Model for Mobile Streaming Videos with Stalling Events," in *Proceedings of SPIE Applications of Digital Image Processing XXXVIII*, San Diego, CA, USA, 2015.
- [3] K. Zeng, H. Yeganeh, and Z. Wang, "Quality-of-experience of Streaming Video: Interactions between Presentation Quality and Playback Stalling," in *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, Phoenix, AZ, USA, 2016.
- [4] International Standards Organization/International Electrotechnical Commission (ISO/IEC), "23009-1:2012 Information Technology – Dynamic Adaptive Streaming over HTTP (DASH) – Part 1: Media Presentation Description and Segment Formats," 2012.
- [5] M. Seufert, T. Höbfeld, and C. Sieber, "Impact of Intermediate Layer on Quality of Experience of HTTP Adaptive Streaming," in *Proceedings of the 11th International Conference on Network and Service Management (CNSM)*, Barcelona, Spain, 2015.
- [6] H. T. Tran, T. Vu, N. P. Ngoc, and T. C. Thang, "A Novel Quality Model for HTTP Adaptive Streaming," in *Proceedings of the 6th IEEE International Conference on Communications and Electronics (ICCE)*, Ha Long, Vietnam, 2016.
- [7] F. Wang, Z. Fei, J. Wang, Y. Liu, and Z. Wu, "HAS QoE Prediction Based on Dynamic Video Features with Data Mining in LTE Network," *Science China Information Sciences*, vol. 60, no. 4, p. 042404, 2017.
- [8] T. Höbfeld, L. Skorin-Kapov, P. Heegaard, and M. Varela, "Definition of QoE Fairness in Shared Systems," *IEEE Communications Letters*, vol. 21, no. 1, pp. 184–187, 2017.
- [9] S. Petrangeli, J. Famaey, M. Claeys, S. Latré, and F. De Turck, "QoE-driven Rate Adaptation Heuristic for Fair Adaptive Video Streaming," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 12, no. 2, pp. 28:1–28:24, 2016.
- [10] J. Chen, M. Ammar, M. Fayed, and R. Fonseca, "Client-Driven Network-level QoE Fairness for Encrypted 'DASH-S'," in *Proceedings of the ACM SIGCOMM Workshop on QoE-based Analysis and Management of Data Communication Networks (Internet-QoE)*, Florianopolis, Brazil, 2016.
- [11] E. Thomas, M. O. van Deventer, T. Stockhammer, A. C. Begen, M.-L. Champel, and O. Oyman, "Applications and Deployments of Server and Network Assisted DASH," in *Proceedings of the International Broadcasting Convention (IBC)*, Amsterdam, The Netherlands, 2016.
- [12] International Standards Organization/International Electrotechnical Commission (ISO/IEC), "23009-5:2015 Information Technology – Dynamic Adaptive Streaming over HTTP (DASH) – Part 5: Server and network assisted DASH (SAND)," 2016.
- [13] International Telecommunication Union, "ITU-T Recommendation P.1203: Parametric Bitstream-based Quality Assessment of Progressive Download and Adaptive Audiovisual Streaming Services over Reliable Transport," 2016. [Online]. Available: <https://www.itu.int/rec/T-REC-P.1203/en>
- [14] —, "ITU-T Recommendation P.910: Subjective Video Quality Assessment Methods for Multimedia Applications," 2008.
- [15] 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," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 4, pp. 187–198, 2015.
- [16] L. De Cicco, V. Calderalo, V. Palmisano, and S. Mascolo, "Elastic: A Client-side Controller for Dynamic Adaptive Streaming over HTTP (DASH)," in *Proceedings of the 20th International Packet Video Workshop (PV)*, San Jose, CA, USA, 2013.
- [17] C. Müller, S. Lederer, and C. Timmerer, "An Evaluation of Dynamic Adaptive Streaming over HTTP in Vehicular Environments," in *Proceedings of the 4th Workshop on Mobile Video (MoVID)*, Chapel Hill, NC, USA, 2012.
- [18] K. Miller, E. Quacchio, G. Gennari, and A. Wolisz, "Adaptation Algorithm for Adaptive Streaming over HTTP," in *Proceedings of the 19th International Packet Video Workshop (PV)*, Munich, Germany, 2012.
- [19] H. Ott, K. Miller, and A. Wolisz, "Simulation Framework for HTTP-Based Adaptive Streaming Applications," in *Proceedings of the Workshop on ns-3 (WNS3)*, Porto, Portugal, 2017.
- [20] 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.
- [21] 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 (CoNEXT)*, Nice, France, 2012.
- [22] 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.
- [23] R. Jain, D.-M. Chiu, and W. R. Hawe, *A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer System*. Eastern Research Laboratory, Digital Equipment Corporation, 1984.
- [24] R. Houdaille and S. Gouache, "Shaping HTTP Adaptive Streams for a Better User Experience," in *Proceedings of the 3rd Annual ACM Conference on Multimedia Systems (MMSys)*, Chapel Hill, NC, USA, 2012.
- [25] F. Wamser, A. Blenk, M. Seufert, T. Zinner, W. Kellerer, and P. Tran-Gia, "Modeling and Performance Analysis of Application-Aware Resource Management," *International Journal of Network Management*, vol. 25, no. 4, pp. 223–241, 2015.
- [26] A. Bentalb, A. C. Begen, and R. Zimmermann, "SDNDASH: Improving QoE of HTTP Adaptive Streaming Using Software Defined Networking," in *Proceedings of the ACM Multimedia Conference (MM)*, Amsterdam, The Netherlands, 2016.
- [27] Y. Li, P. A. Frangoudis, Y. Hadjadj-Aoul, and P. Bertin, "A Mobile Edge Computing-based Architecture for Improved Adaptive HTTP Video Delivery," in *Proceedings of the IEEE Conference on Standards for Communications and Networking (CSCN)*, Berlin, Germany, 2016.
- [28] G. Cofano, L. De Cicco, T. Zinner, A. Nguyen-Ngoc, P. Tran-Gia, and S. Mascolo, "Design and Experimental Evaluation of Network-assisted Strategies for HTTP Adaptive Streaming," in *Proceedings of the 7th ACM International Conference on Multimedia Systems (MMSys)*, Klagenfurt, Austria, 2016.
- [29] J. Postel, "Transmission Control Protocol," Internet Requests for Comments, RFC 793, 1981. [Online]. Available: <https://tools.ietf.org/html/rfc793>
- [30] M. Allman, V. Paxson, and E. Blanton, "TCP Congestion Control," Internet Requests for Comments, RFC 5681, 2009. [Online]. Available: <https://tools.ietf.org/html/rfc5681>
- [31] T.-Y. Huang, N. Handigol, B. Heller, N. McKeown, and R. Johari, "Confused, Timid, and Unstable: Picking a Video Streaming Rate is Hard," in *Proceedings of the Internet Measurement Conference (IMC)*, Boston, MA, USA, 2012.
- [32] Blender Foundation, "Big Buck Bunny," 2008. [Online]. Available: <https://www.bigbuckbunny.org>
- [33] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image Quality Assessment: From Error Visibility to Structural Similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.