




Article

Cross-Protocol Unfairness between Adaptive Streaming Clients over HTTP/3 and HTTP/2: A Root-Cause Analysis

Chanh Minh Tran ^{1,*} , Tho Nguyen Duc ¹, Phan Xuan Tan ^{2,*}  and Eiji Kamioka ^{1,*} ¹ Graduate School of Engineering and Science, Shibaura Institute of Technology, Tokyo 135-8548, Japan; nb20501@shibaura-it.ac.jp² Department of Information and Communications Engineering, Shibaura Institute of Technology, Tokyo 135-8548, Japan

* Correspondence: nb20502@shibaura-it.ac.jp (C.M.T.); tanpx@shibaura-it.ac.jp (P.X.T.); kamioka@shibaura-it.ac.jp (E.K.)

Abstract: With the introduction of HTTP/3, whose transport is no longer the traditional TCP protocol but the novel QUIC protocol, research for solutions to the unfairness of Adaptive Streaming over HTTP (HAS) has become more challenging. In other words, because of different transport layers, the HTTP/3 may not be available for some networks and the clients have to use HTTP/2 for their HAS applications instead. Therefore, the scenario in which HAS over HTTP/3 (HAS/3) competes against HTTP/2 (HAS/2) must be considered seriously. However, there has been a shortage of investigations on the performance and the origin of the unfairness in such a cross-protocol scenario in order to produce proper solutions. Therefore, this paper provides a performance evaluation and root-cause analysis of the cross-protocol unfairness between HAS/3 and HAS/2. It is concluded that, due to differences in the congestion control mechanisms of QUIC and TCP, HAS/3 clients obtain larger congestion windows, thus requesting higher video bitrates than HAS/2. As the problem lies in the transport layer, existing client-side ABR-based solutions for the unfairness from the application layer may perform suboptimally for the cross-protocol case.

Keywords: adaptive streaming; HTTP/3; QUIC; cross-protocol; unfairness; congestion control

Citation: Tran, C.M.; Nguyen Duc, T.; Tan, P.X.; Kamioka, E. Cross-Protocol Unfairness between Adaptive Streaming Clients over HTTP/3 and HTTP/2: A Root-Cause Analysis. *Electronics* **2021**, *10*, 1755. <https://doi.org/10.3390/electronics10151755>

Academic Editors: Manuel Perez Malumbres, Carlos Tavares Calafate and Glenn Van Wallendael

Received: 7 June 2021

Accepted: 20 July 2021

Published: 21 July 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Adaptive streaming over HTTP (HAS) has become the de facto standard for most of the online video services on the internet nowadays, thanks to its ability to instantaneously adapt the video quality with the network condition [1]. In HAS, a video at the streaming server is encoded into multiple quality versions (in terms of bitrates), each of which is split into multiple small segments with the same duration. Then, every streaming client is equipped with an Adaptive Bitrate Algorithm (ABR) in its video player to continuously monitor the network condition, therefore requesting the suitable bitrate for every video segment. Accordingly, undesirable incidents such as playback stalls and quality variation can be reduced, thus maintaining a high quality of experience (QoE) for the users.

Nevertheless, in order to cope with the constantly increasing number of online video users [2], as well as their daily usage time of the service [3], the performance of the HAS services against multiple competing clients still requires further optimization. In multi-client scenarios, it has been proven that due to the mismatch of the clients' downloading states originating from their ABRs on the application layer, some clients may overestimate their bandwidths and request higher video bitrates than others [4–6]. Such a phenomenon is defined as the unfairness problem, which causes QoE deterioration and a negative impact on the user retention rate. Therefore, over the years, various research has been conducted and various solutions to the unfairness have been proposed [7–10].

On the other hand, the protocol stacks have seen some significant changes in recent years that complicate the efforts to solve the unfairness of HAS. In fact, the HTTP/3

protocol [11] was proposed and is expected to be standardized in the near future. While HTTP/2 differs from HTTP/1.1 by the novel features on the application layer [12], the major difference between HTTP/3 and its successors lies in the transport protocol. In other words, while HTTP/1.1 and HTTP/2 have been running on top of the well-tuned TCP protocol for decades, the HTTP/3 utilizes the novel QUIC protocol [13] for its transport. QUIC actually runs atop the UDP protocol, which is often blocked or limited by network entities due to known security risks [14–16]. This means that, as QUIC relies on UDP, there are realistic scenarios that a client fails to use HTTP/3 because of configurations of its network system and has to use HTTP/2 or HTTP/1.1 instead (Figure 1). Thus, streaming providers must support both HTTP/3 and the former HTTP versions at the same time to ensure service availability. This raises the need for the performance investigation of the unfairness with regard to the cross-protocol scenario, where streaming clients using HTTP/3 concurrently compete against ones using HTTP/2 or HTTP/1.1.

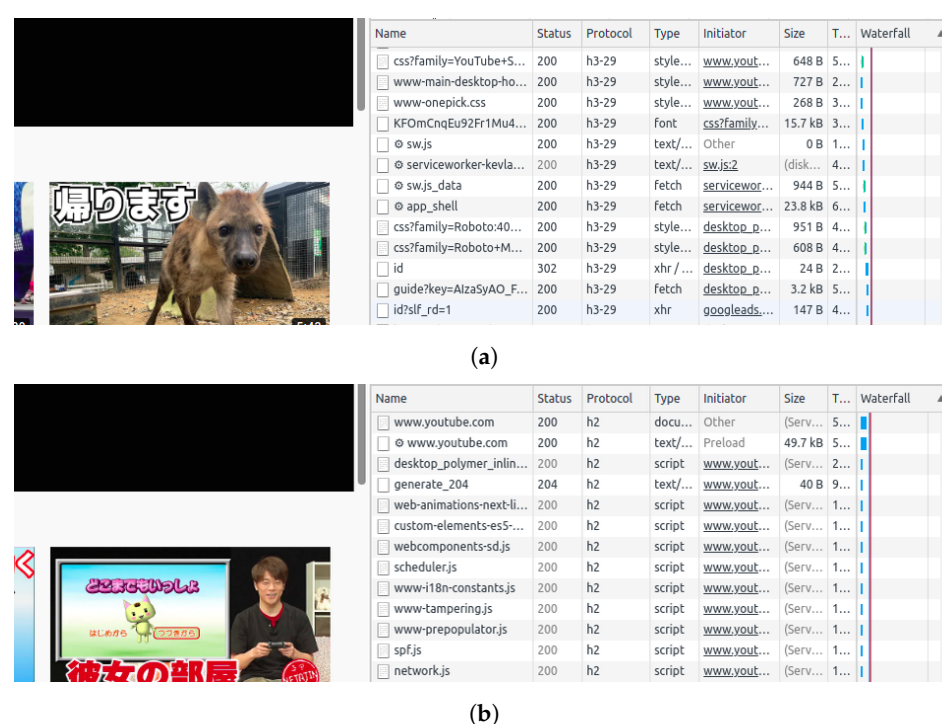


Figure 1. An example realistic case of Youtube where the HTTP/3 cannot be utilized. (a) A normal client that successfully sends HTTP/3 requests (h3-29); (b) a firewall-controlled client that fails to send HTTP/3 requests and has to use HTTP/2 (h2) instead.

In this manner, research about the cross-protocol unfairness between HAS over HTTP/3 (HAS/3) and over HTTP/2 (HAS/2) or HTTP/1.1 (HAS/1.1) should attract more attention. However, to the best of our knowledge, only a few works have attempted to investigate the unfairness in such a case for HAS and no solution has ever been proposed due to the following limitations. Firstly, as HTTP/3 and QUIC are still under standardization, their features and characteristics vary upon drafted versions and their implementations vary upon releases. Consequently, performance conclusions in the existing works are contradictory against one another. Secondly, those works only provide observations and lack root-cause analysis. Without such crucial information, research for an effective solution to the cross-protocol unfairness could not be conducted.

Realizing the aforementioned drawbacks, this paper presents an up-to-date performance evaluation and a comprehensive root-cause investigation on the cross-protocol unfairness of HAS/3 against its successors HAS/2 and HAS/1.1. In fact, the HTTP/2 nowadays accounts for 64% of HTTP requests over the internet and is expected to continue growing linearly [17]. With such usage growth, in addition to the more advanced features that the HTTP/2 provides [12], it is believed that HTTP/1.1 will become degraded

in the near future. Therefore, this work focuses mainly on the competition of HAS/3 and HAS/2 clients. Our analysis demonstrates that, with the newest documentation and implementation, HTTP/3 clients tend to experience higher video quality than HTTP/2 clients when they compete under the same bottleneck network. Taking an in-depth look at their transport—QUIC for HTTP/3 versus TCP for HTTP/2—it is found that, despite both implementations running the same congestion control algorithm, QUIC is able to acquire a larger congestion window than TCP. Such a behavior actually originates from the different mechanisms of the congestion control and the characteristics of QUIC and TCP [18]. In detail, we argue that the differences in loss epoch life cycle, acknowledgement (ACK) ranges and minimum congestion window size, as well as the unreliable nature of the UDP (QUIC's based protocol), allow QUIC to receive ACK frames more frequently than TCP. As a consequence, QUIC updates its congestion window more aggressively than TCP, resulting in a higher occupation of the bandwidth. Moreover, since such an under-performance is transparent to the application layer, it has been proven that the existing unfairness solutions utilizing the client-side ABR-based approach on the application layer may fail to achieve desirable effectiveness. Instead, follow-up research should focus on either examining server- or network-based approaches on the transport layer, or tweaking the functionalities and parameters of the transport QUIC. In summary, the distinguished contributions of this work are as follows:

- An up-to-date investigation of the cross-protocol unfairness between HAS/3 and HAS/2 is provided, showing that HAS/3 clients unfairly experience higher video quality.
- A comprehensive analysis of such an unfairness problem is conducted, which concludes that its origin lies in the differences in congestion control mechanisms and the characteristics of the transport layer QUIC for HTTP/3 and TCP.
- Based on the root-cause analysis, we provide suggestions for proper solutions to the cross-protocol unfairness between HAS/3 and HAS/2.

The remainder of this paper is organized as follows: Section 2 reviews the existing research about the cross-protocol unfairness between HAS/3 and HAS/2. The hypothesis on the cause of the cross-protocol unfairness and the methodology and experimental setup for validation are presented in Section 3. The experimental results are analyzed and discussed in Sections 4 and 5, respectively. Finally, Section 6 concludes this paper.

2. Related Work

Ever since the proposal of the QUIC transport protocol and the HTTP/3, a dedicated track of research has focused on evaluating its influences on the performance of the HAS. To name a few, in [19], the authors showed that HTTP-over-QUIC, which was the former name of HTTP/3 before November 2018 [20], helped the client start the video quicker, reduced the rebuffering ratio and improved the average bitrate compared to HTTP-over-TCP in WiFi, LTE and 3G networks. Moreover, for the situation when the client switched between network interfaces, the results also expressed the outperformance of HTTP-over-QUIC in those metrics. The work in [21] combined a traditional video streaming approach and a retransmission technique for Scalable High Efficiency Video Coding streaming over HTTP/3. Compared with HTTP/2, the proposed approach provided higher average video quality and more stable bitrate switches when packet loss and retransmission occurred. Meanwhile, the paper [22] discussed and provided suitable settings of the QUIC's parameters for better performance of the HAS. On the other hand, [23] investigated the impact of HTTP/3 and QUIC on 360° HAS and argued that they could effectively reduce rebuffering counts and duration. Despite the fact that these works provided promising performances of the HTTP/3 and QUIC on HAS, they did not consider the multi-client scenario, where the unfairness problem requires more attention.

As explained in the previous section, the differences in the transport protocol between HTTP/3 and HTTP/2 or HTTP/1.1 highlight the need for tackling the cross-protocol unfairness of the clients requesting these traffics. However, research about such a problem for HAS is surprisingly limited. The work in [24] provided a performance study of several

ABRs under the HTTP-over-QUIC. The author did not state which version of QUIC had been tested in their work. However, they did mention the Google QUIC, which is the very first variant of today's IETF QUIC. Based on their evaluation results, it was shown that the streaming clients using HTTP-over-QUIC failed to experience a video bitrate as high as that of the HTTP-over-TCP clients. The author argued that such an underperformance of the HTTP-over-QUIC clients was because most ABRs were tuned to work well with TCP so they did not utilize the features of QUIC effectively. In their later work [25], the QUIC retransmission scheme was employed with Google QUIC version Q043, showing promising improvements in video quality and bitrate stability. Interestingly, while they noted that their solution might result in unfairness among the clients running HTTP-over-QUIC, they also observed that such a client also starved out the HTTP-over-TCP clients. Such an observation was contradictory to their own previous work in [24]. On the other hand, a performance investigation similar to that of [24] was conducted in [26] with the Google QUIC server v39. The finding in [26] showed that, while the HAS clients running over QUIC could perform fairly against one another, they provided 37% higher bitrate than the HAS over TCP clients in most cross-protocol cases.

It can be noticed that the findings in these existing works were contradictory to one another and were actually outdated as the Google QUIC is now deprecated to make way for the IETF QUIC [27]. Moreover, their works only provided observation and lacked in-depth analysis on the origin of the problem. Thus, to the best of our knowledge, there have been no solutions proposed for the cross-protocol unfairness between HAS/3 and HAS/2 clients due to such a serious drawback. In order to clarify the root cause of the cross-protocol unfairness, in the following sections, we present our hypothesis, experimental methodology and performance analysis with the latest version of the HTTP/3 and QUIC.

3. Hypothesis and Methodology

This section describes the hypothesis on the existence and the cause of the cross-protocol unfairness among HAS/3 and HAS/2 clients in detail, as well as the methodology and experimental settings for validating it.

3.1. Hypothesis

In HAS, a client's ABR typically decides the video bitrate based on its observation of the available bandwidth. Apparently, unfair bandwidth consumption of clients towards one another can lead to an unfair bitrate selection. Nevertheless, the bandwidth fairness is actually well-handled by the congestion control of the TCP protocol in the transport layer [28,29] and existing research also concluded that the unfairness in bitrate selection of single-protocol HAS/2 or HAS/1.1 was the result of inefficient ABRs in the application layer [4]. Therefore, before the introduction of HTTP/3, previous investigations and solutions to the unfairness only focused on optimizing the ABR from the client side.

On the other hand, this work considers the cross-protocol scenario, where TCP is no longer the only transport layer of all clients. In fact, the transport QUIC of HAS/3 also implements TCP's congestion control. However, as stated in the draft RFC, there exist several differences, which are considered as enhancements, in the mechanism of QUIC's congestion control compared with TCP's [18]. The differences that are most relevant to this paper are as follows:

- Immediate reaction to packet loss: When a packet loss is detected, QUIC typically reacts every round trip. For the case of TCP, it may have to wait for multiple round trips before progressing with follow-up actions.
- More ACK ranges for loss recovery: While TCP utilizes only three selective ACKs for loss recovery, QUIC supports many ACK ranges to speed up such a process, especially in high-loss environments.
- Increasing minimum congestion window: QUIC recommends that the minimum congestion window is two packets long, while that of the TCP is only one packet.

Thus, it is likely that the performance of their congestion control is dissimilar and causes uneven bandwidth distribution. As the clients' ABRs rely heavily on the bandwidth obtained by their transport, their bitrate selections are affected, which causes unfairness in video bitrates. To clarify this hypothesis, the experimental methodology and settings of this paper are presented in the remaining parts of this section.

3.2. Methodology

3.2.1. Experimental Scenarios

In order to confirm the cross-protocol unfairness between HAS/3 and HAS/2 with the latest documentation and implementation of HTTP/3, as well as to validate our hypothesis on the root cause of the problem, an experimental evaluation is conducted. The experiment is run with several total bandwidth limits B to assess the performance of the clients across different competitiveness levels of the network (i.e., a lower value of B means a more serious competition). In this experiment, the clients are examined with $B \in \{3, 5\}$ (mbps).

In addition, the number of each type of client (i.e., HAS/3 or HAS/2 clients) also varies. This is to verify whether the unfairness is explicitly towards a particular type of client. In this experiment, the total number of all clients is varied from 2 to 3 clients. For the case of 2 clients, there is only the scenario that 1 HAS/3 client competes with 1 HAS/2 client. Meanwhile, with 3 clients, there are two subscenarios to be considered: 1 HAS/3 client competes with 2 HAS/2 clients, and 2 HAS/3 clients compete with 1 HAS/2 client. In summary, the details of the evaluated scenarios are shown in Table 1.

Table 1. Summary of the experimental scenarios.

Total Number of Clients	Scenario	Denotation
2 clients	1 HAS/3 client and 1 HAS/2 client	S_{11}
3 clients	1 HAS/3 client and 2 HAS/2 clients	S_{12}
	2 HAS/3 clients and 1 HAS/2 client	S_{21}

3.2.2. Investigation Metrics

Existing works on the unfairness of HAS often utilized the average unfairness index \bar{F} of the video bitrates selected by the clients' ABRs as a numerical metric of the unfairness level [6–8,10]. Therefore, in this paper, this metric will be considered to investigate such a problem. \bar{F} , adopted from [7], is estimated by averaging the unfairness index F from the beginning to the end of a streaming session, with F calculated based on the discrimination index of the Jain Fairness [30] as in Equation (1).

$$F = \sqrt{1 - \text{JainFair}} = \sqrt{1 - \frac{(\sum_{c=1}^C r_{c,t})^2}{C * \sum_{c=1}^C r_{c,t}^2}} \quad (1)$$

where $r_{c,t}$ denotes the bitrate r selected by the client c at time t . A larger value of \bar{F} indicates a more severe unfairness. Additionally, the average bitrate \bar{r} of each client throughout its streaming session will be measured in order to judge whether higher bitrates are only available at a specific type of client.

Furthermore, in this work, it is hypothesized that the differences in congestion control of the transport protocol QUIC and TCP actually cause the unfairness in bandwidth consumption among the clients. In order to clarify this hypothesis, the time-varying congestion window of each client, which is the output of the congestion control algorithm, is also collected.

3.3. Experimental Setup

In this subsection, the description of the experimental settings for investigating the cross-protocol unfairness between HAS/3 and HAS/2 is provided. Figure 2 depicts the experiment topology.

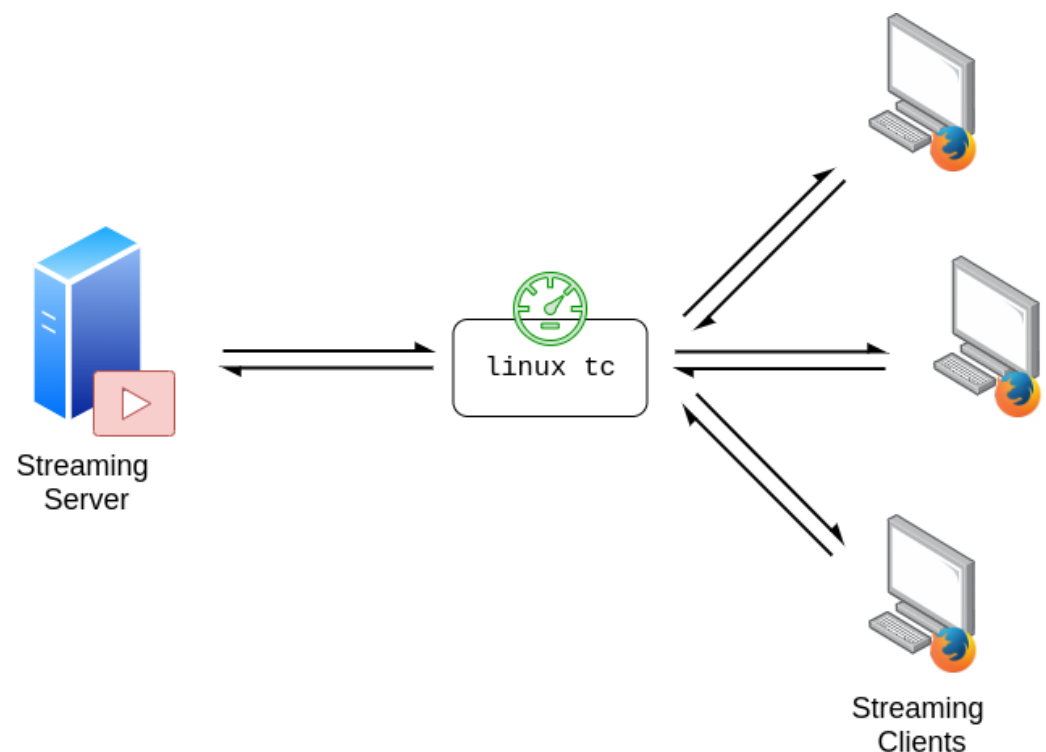


Figure 2. Experimental topology.

The video server and clients were deployed on separate virtual machines running Ubuntu 20.04 LTS with 4GB of RAM and 4 processor cores. These machines were actually virtualized on a Core i5 physical machine running Ubuntu 20.04 LTS with 80 GB of RAM. For enabling HTTP/3, the server applied the implementation of *quic-go* v0.20.1 [31], which supported version 34 of the drafted HTTP/3 and QUIC, which was the latest draft at the time we conducted this experiment. As for HTTP/2, the server simply used the native *http* package of Golang [32]. Meanwhile, the total bandwidth limit B was configured using the *linux tc* [33]; the congestion window was extracted from the *qlog* files of *quic-go* for HTTP/3 and *linux ss* [34] for HTTP/2.

The streaming application was packaged at the server based on the *dash.js* framework [35], which was run via a Firefox web browser from the client side. The server provided 300 2-second segments of the open-source *Big Buck Bunny* video with 11 quality versions, i.e., {100, 200, 300, 500, 700, 900, 1200, 1500, 2000, 2500, 3000} (kbps). At the client, the maximum buffer was set to 30 seconds. In this experiment, the clients' performances were tested with two types of ABR: general ABR and fairness ABR. For the general ABR, the default ABR of *dash.js* was employed, which was actually a dynamic ABR that conditionally switches between its own throughput ABR and the well-known buffer-based BOLA [36]. The *dash.js* ABR was not tuned for solving the unfairness issues, whereas we were also curious about whether existing fairness ABRs could perform well in such a cross-protocol scenario. For this reason, the FESTIVE [7], which was among the most famous baseline fairness ABRs (i.e., ABRs taking into account the fairness optimization in their designed goals), was considered. The FESTIVE utilized the harmonic mean of the estimated bandwidth, gradual and stateful bitrate transition, and randomized segment download scheduler to overcome the unfairness from the client side.

The experiment of each scenario described in Section 3.2 was tested 10 times to ensure performance consistency. In the following section, the detailed results and analysis are presented.

4. Results and Analysis

This section provides the results and analysis of our experiment on the cross-protocol unfairness between HAS/3 and HAS/2 clients. For the numerical results (\bar{F} and \bar{r}) of each scenario, we show the average values of all 10 runs. As for the visualized time-varying figures, due to similar behaviors, only one representative run for each scenario is shown.

4.1. General ABR

In this subsection, the results collected for the general ABR are analyzed. Table 2 summarizes the average bitrate \bar{r} of every client. In this table, $C_{HAS/Y}^X$ denotes the client X running HAS/ Y with $Y \in \{2, 3\}$. Meanwhile, the average unfairness index \bar{F} of all scenarios is shown in Table 3.

Table 2. The average bitrate \bar{r} (kbps) of each client running the general ABR.

B	S_{11}		S_{12}			S_{21}		
	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/2}^3$	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/3}^3$
3 mbps	1412 \pm 32	1007 \pm 31	1105 \pm 46	706 \pm 32	717 \pm 27	897 \pm 70	702 \pm 43	918 \pm 41
5 mbps	2325 \pm 33	1552 \pm 45	1684 \pm 105	1187 \pm 107	1180 \pm 96	1536 \pm 106	1115 \pm 124	1483 \pm 87

Table 3. The average unfairness index \bar{F} of the general ABR.

B	S_{11}	S_{12}	S_{21}
3 mbps	0.1895 \pm 0.0195	0.2711 \pm 0.0272	0.2122 \pm 0.0150
5 mbps	0.2105 \pm 0.0119	0.2398 \pm 0.0275	0.2259 \pm 0.0180

The numerical results in Table 3 quantitatively demonstrate that, across all scenarios, the bitrate decisions of same-type HAS clients were relatively indistinguishable. In contrast, it was obvious that the HAS/3 clients ended up requesting higher bitrates than the HAS/2 clients. In particular, the bitrates selected by the HAS/3 clients were 42.05% higher than those of the HAS/2 clients on average. This led to the results of the unfairness index \bar{F} in Table 3. Although such an unfairness performance was actually predictable because the general ABR was not tuned for solving such a problem, the fact that the high bitrates were reserved only for the HAS/3 clients was very interesting. Taking a closer look at this observation, Figures 3 and 4 depict the time-varying bitrates and moving average congestion windows ($cwnd_ma$) of the clients under 3 mbps and 5 mbps, respectively.

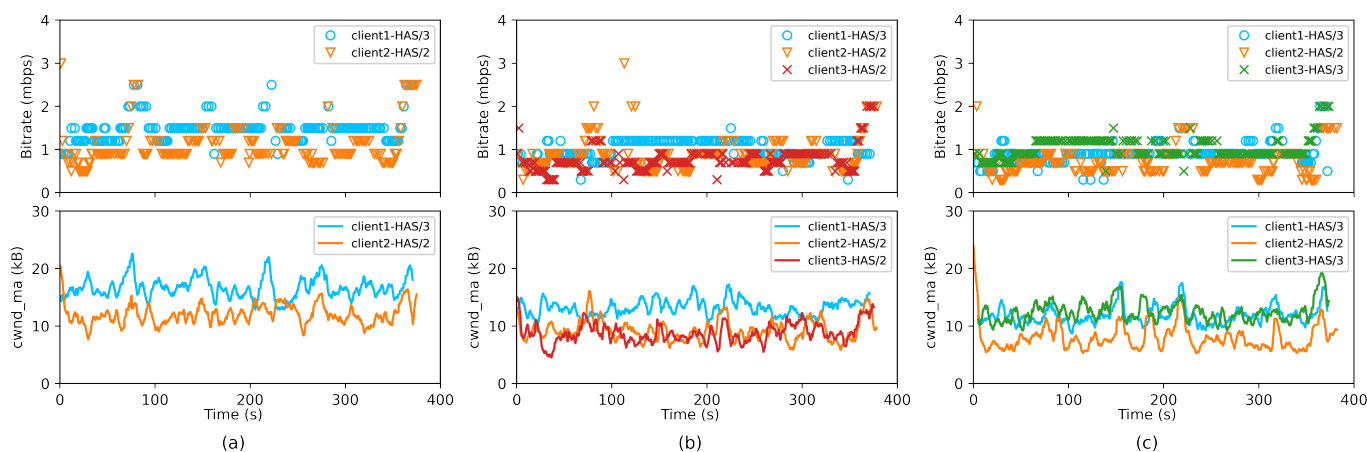


Figure 3. The representative time-varying bitrates and $cwnd_ma$ of the clients running the general ABR with $B = 3$ (mbps), under the scenario (a) S_{11} , (b) S_{12} and (c) S_{21} .

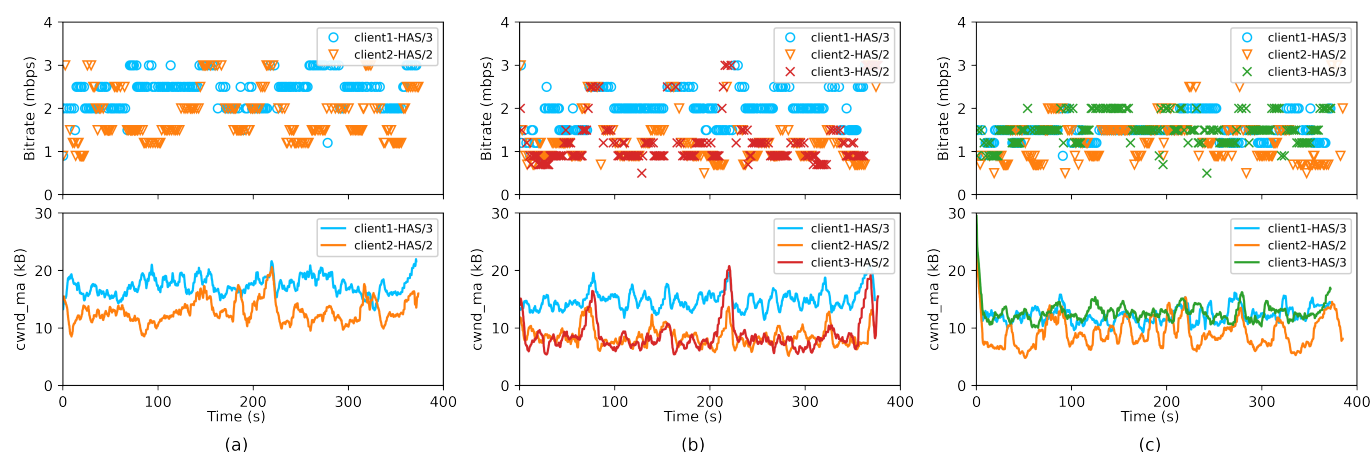


Figure 4. The representative time-varying bitrates and $cwnd_ma$ of the clients running the general ABR with $B = 5$ (mbps), under the scenario (a) S_{11} , (b) S_{12} and (c) S_{21} .

It can be observed from Figures 3 and 4 that the bitrates of both types of clients varied significantly; they were occasionally increased or decreased with large amplitudes and within short periods. This behavior actually harmed the fairness as such abrupt and large changes might significantly increase the difference in clients' bitrates at some time intervals. More importantly, the time-varying $cwnd_ma$ from Figures 3 and 4 expresses that the HAS/3 clients always received a higher congestion window ($cwnd$) than the HAS/2 clients. Meanwhile, the $cwnd_ma$ among the HAS/3 clients or among the HAS/2 clients were relatively comparable. This means that, despite running the same congestion control algorithm, the transport QUIC of HAS/3 clients helped them obtain more $cwnd$, thus receiving a greater portion of the bandwidth than the HAS/2 clients running on TCP. This phenomenon strongly supports our hypothesis in Section 3.1, that the performance of the congestion control between QUIC and TCP was different, which harmed the fairness of their bandwidth consumption. As a result, the HAS/3 clients running on QUIC were able to request higher bitrates. Moreover, since this problem lies in the transport layer, it is highly possible that such a behavior also occurred with the fairness ABR. As the ABR worked on the application layer, it could not control but only relied on the bandwidth given by its transport layer. To validate this expectation, the next subsection shows similar analysis results of the fairness ABR.

4.2. Fairness ABR

Similar to the previous subsection, Table 4 shows the measured \bar{F} , Table 5 summarizes \bar{r} of every client, while Figures 5 and 6 illustrate the time-varying performance under 3 mbps and 5 mbps, respectively, in terms of bitrates and $cwnd_ma$ of the fairness ABR.

Table 4. The average unfairness index \bar{F} of the fairness ABR.

B	S_{11}	S_{12}	S_{21}
3 mbps	0.0814 ± 0.0262	0.1095 ± 0.0227	0.0991 ± 0.0245
5 mbps	0.0697 ± 0.0129	0.0909 ± 0.0106	0.0712 ± 0.0120

Table 5. The average bitrate \bar{r} (kbps) of each client running the fairness ABR.

B	S_{11}		S_{12}			S_{21}		
	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/2}^3$	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/3}^3$
3 mbps	1106 ± 34	927 ± 41	866 ± 30	706 ± 19	689 ± 30	804 ± 24	653 ± 19	792 ± 36
5 mbps	1720 ± 21	1461 ± 41	1326 ± 18	1071 ± 21	1138 ± 16	1234 ± 25	1064 ± 39	1200 ± 29

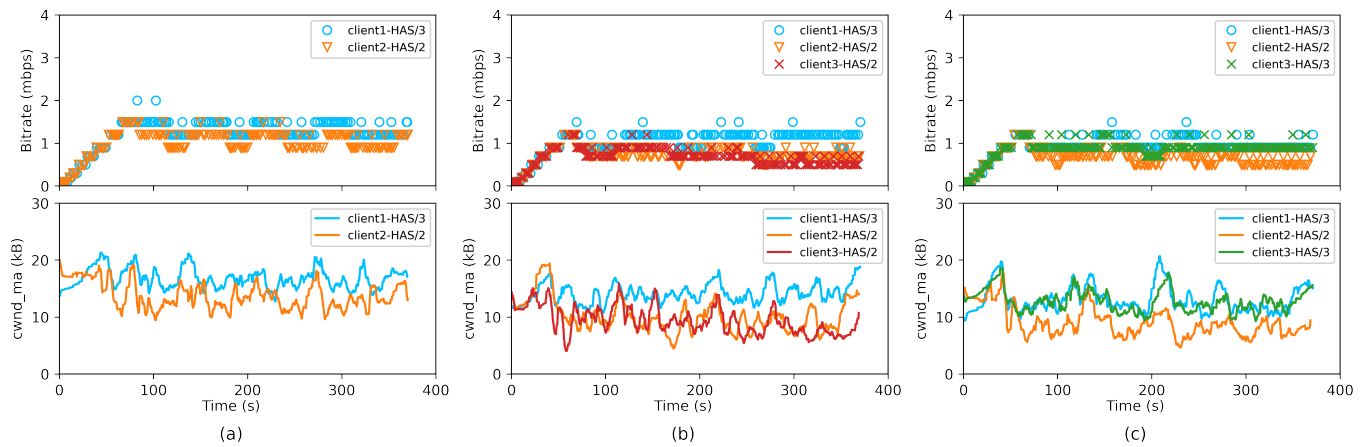


Figure 5. The representative time-varying bitrates and *cwnd_ma* of the clients running the fairness ABR with $B = 3$ (mbps), under the scenario (a) S_{11} , (b) S_{12} and (c) S_{21} .

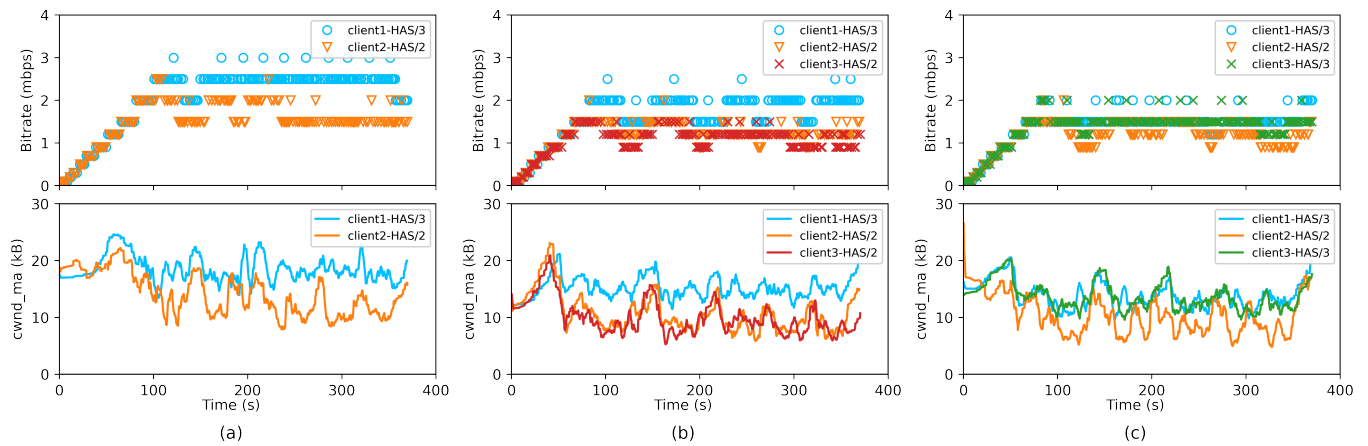


Figure 6. The representative time-varying bitrates and *cwnd_ma* of the clients running the fairness ABR with $B = 5$ (mbps), under the scenario (a) S_{11} , (b) S_{12} and (c) S_{21} .

It can be inferred from Table 4 that the conditions of the cross-protocol unfairness across all scenarios were significantly better than the general ABR. Indeed, the fairness ABR provided an average of 44.69% improvement in \bar{F} compared to the general ABR. However, Table 5 concludes a similar tendency to that described in Section 4.1: the HAS/3 clients still experienced higher video bitrates than HAS/2 clients, despite the fact that the gap between them was minimized to 19.47% higher for HAS/3 on average. Observing Figures 5 and 6, we can see that the clients' bitrates were more in control and tended to vary within specific ranges. This behavior was the result of the gradual and stateful update mechanism of the fairness ABR: the client stayed at every bitrate level for a specific number of segments before, if necessary, increasing or decreasing it to the closest higher or lower level. Thus, there were no abruptly large bitrate changes and the difference in the clients' average bitrates was smaller. Moreover, their bitrate selections were relatively identical for a specific amount of time at the beginning. This was because the fairness ABR forced all clients to choose the lowest bitrate at the beginning of their streaming sessions to avoid playback stalls. Such a strategy not only explains why the results of \bar{r} of all types of client with all total bandwidth limits were smaller when running the fairness ABR, but also contributes to the minimization of the bitrate difference discussed earlier. Nevertheless, as expected, the HAS/3 clients were given higher *cwnd*, thus requesting higher bitrates due to the higher occupation of the bandwidth. As a result, although the bitrate selections were similar at the beginning due to the ABR's strategy, they ended up varying unfairly based on their estimations of bandwidth afterwards.

In summary, the results of both types of ABR demonstrate that, under a cross-protocol scenario, the bitrate selection among the clients was always unfairly higher for the HAS/3 clients. This is because the transport QUIC allowed its HAS/3 clients to utilize higher *cwnd* than the TCP of the HAS/2 clients. Therefore, the root cause of such cross-protocol unfairness lies in the congestion control mechanisms of QUIC and TCP, rather than in the application layer, as argued in previous investigations that only focused on the single-protocol unfairness.

5. Discussion

Based on the analysis in Section 4, it is found that the transport QUIC was able to achieve higher *cwnd* for its HAS/3 clients than TCP for HAS/2, leading to an unfair distribution of the shared bandwidth. For this reason, even the referenced fairness ABR could not perform optimally to provide similar bitrates for both types of HAS client. Such a phenomenon is actually explainable by examining the enhancements of QUIC's congestion control in comparison with TCP's, as described in Section 3.2. These enhancements actually empower QUIC to transmit data packets much faster than TCP. In addition, QUIC runs atop UDP, which is naturally faster than TCP due to the unreliable characteristics [16]. For these reasons, QUIC is able to receive ACK frames faster than TCP, so that QUIC updates its *cwnd* more aggressively and occupies more than the fair share of bandwidth. In order to confirm this explanation, Tables 6 and 7 show the average number of *cwnd* updates of the clients using the general and the fairness ABR, respectively, in the previous experiment; meanwhile, Figure 7 illustrates an example of the zoomed-in time-varying *cwnd* of the clients using the fairness ABR under the scenario S_{11} with $B = 3$ (mbps).

Table 6. The average number of *cwnd* updates of each client running the general ABR.

B	S_{11}		S_{12}			S_{21}		
	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/2}^3$	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/3}^3$
3 mbps	4407 ± 238	2586 ± 351	3776 ± 287	2569 ± 332	2610 ± 276	3894 ± 255	2611 ± 363	3971 ± 253
5 mbps	6916 ± 349	3631 ± 402	6662 ± 387	3637 ± 458	3600 ± 484	6250 ± 390	3851 ± 472	6048 ± 399

Table 7. The average number of *cwnd* updates of each client running the fairness ABR.

B	S_{11}		S_{12}			S_{21}		
	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/2}^3$	$C_{HAS/3}^1$	$C_{HAS/2}^2$	$C_{HAS/3}^3$
3 mbps	3443 ± 198	2212 ± 235	3425 ± 185	2076 ± 199	2050 ± 203	3397 ± 188	2054 ± 221	3301 ± 197
5 mbps	5360 ± 232	3159 ± 221	5001 ± 243	3091 ± 233	2969 ± 290	4836 ± 283	3110 ± 211	4815 ± 201

Obviously, from Tables 6 and 7, the HAS/3 clients updated the *cwnd* much more frequently than the HAS/2 clients. The illustration in Figure 7 also supports this conclusion. It can be noticed that the number of *cwnd* updates of all clients running the general ABR was higher than those running the fairness ABR across all scenarios and bandwidth limits. This was simply because the clients running the general ABR often abruptly requested bitrates much higher than the fair portion of the total bandwidth limit, leading to higher packet loss due to insufficient bandwidth and increasing the frequency of *cwnd* updates.

In conclusion, our hypothesis on the root cause of the cross-protocol unfairness of HAS/3 and HAS/2 was verified; in other words, the dissimilarities in congestion control of QUIC and TCP led to unfair bandwidth allocation and, finally, unfair bitrate selections. As the the problem arises from the transport layer, follow-up research on the cross-protocol unfairness can investigate existing methods working on this layer, such as server- or network-based solutions. For example, [10,37] utilized a bandwidth allocation module that assigned an equal and separate bandwidth slice for each client. Figure 8 visualizes the bitrate selection performance of the scenario S_{11} with $B = 3$ (mbps) when applying such a method.




Figure 7. An example zoomed-in time-varying *cwnd* of the clients using the fairness ABR under the scenario S_{11} with $B = 3$ (mbps).




Figure 8. An example time-varying bitrate selection of the scenario S_{11} with $B = 3$ (mbps) when applying a bandwidth allocation method.

The time-varying bitrate selection clearly demonstrates superior fairness performance compared with the client-side fairness ABR tested in Section 4.2. This is because, as explained in [10,37], when every client had its own specific bandwidth, they ultimately did not compete with one another and only maximized the bitrates based on the assigned bandwidth. Consequently, since all clients were given an equal bandwidth, their bitrate selections ended up being fair regardless of their HTTP versions or transport protocols. Nevertheless, the server- and network-based solutions have been questioned regarding their consistent efficiency in large-scale networks due to the extra computational complexity and overhead, or regarding the deployment feasibility as they require additional network entities [38,39]. The cost benefit regarding these matters should be carefully considered before applying such solutions in real life.

On the other hand, the transport QUIC is actually implemented on the user space of both endpoints [40]. Therefore, modifications of the protocol's functionalities and parameters become more feasible as they do not experience the ossification problem caused

by conservative network blocks on-the-fly [41,42]. For this reason, future research can also consider tweaking the functionalities and parameters of QUIC for obtaining a fairer bandwidth for the HAS/3 clients against HAS/2 and/or HAS/1.1.

6. Conclusions and Future Work

This paper presents an up-to-date performance analysis of the HAS/3 and HAS/2 clients under a cross-protocol scenario and confirms that HAS/3 clients always unfairly acquire higher bitrates than its successors. Looking into the transport layer, it is found that the root cause of this underperformance lies in the aggressive occupation of the congestion window of QUIC—the transport of HAS/3. Such a behavior originates from enhancements in the congestion control mechanism of QUIC, as well as its different characteristics with TCP. As a result, it is proven that the existing client-side fairness ABR fails to provide fair bitrates for the clients due to irrelevant working layers. Based on this conclusion, for future work, we will focus on tuning the congestion control algorithm of QUIC. In addition, the feasibility of server- and network-based unfairness solutions will also be examined under such a cross-protocol scenario.

Author Contributions: Conceptualization, C.M.T., T.N.D. and P.X.T.; Methodology, C.M.T., T.N.D. and P.X.T.; Supervision, P.X.T. and E.K.; Writing—original draft, C.M.T. and T.N.D.; Writing—review and editing, P.X.T. and E.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Stockhammer, T. Dynamic Adaptive Streaming over HTTP—: Standards and Design Principles. In Proceedings of the Second Annual ACM Conference on Multimedia Systems, MMSys '11, San Jose, CA, USA, 23 February 2011; Association for Computing Machinery: New York, NY, USA, 2011; pp. 133–144. [\[CrossRef\]](#)
2. *Global Digital Video 2019*; Technical Report; eMarketer: New York, NY, USA, 2019.
3. *Online Video Forecasts 2019*; Technical Report; Zenith: New York, NY, USA, 2019.
4. Akhshabi, S.; Begen, A.C.; Dovrolis, C. An Experimental Evaluation of Rate-Adaptation Algorithms in Adaptive Streaming over HTTP. In Proceedings of the Second Annual ACM Conference on Multimedia Systems, MMSys '11, San Jose, CA, USA, 23 February 2011; Association for Computing Machinery: New York, NY, USA, 2011; pp. 157–168. [\[CrossRef\]](#)
5. Akhshabi, S.; Anantakrishnan, L.; Begen, A.C.; Dovrolis, C. What Happens When HTTP Adaptive Streaming Players Compete for Bandwidth? In Proceedings of the 22nd International Workshop on Network and Operating System Support for Digital Audio and Video, NOSSDAV '12, Amherst, MA, USA, 21 June 2012; Association for Computing Machinery: New York, NY, USA, 2012; pp. 9–14. [\[CrossRef\]](#)
6. Wang, Y.; Tran, C.M.; Duc, T.N.; Wu, X.; Tan, P.X.; Kamioka, E. An Experimental Study on The Unfairness in Adaptive Streaming with HTTP/2 Server Push. In Proceedings of the 2019 International Conference on Video, Signal and Image Processing, VSIP 2019, Wuhan, China, 29–31 October 2019; Association for Computing Machinery: New York, NY, USA, 2019; pp. 94–98. [\[CrossRef\]](#)
7. Jiang, J.; Sekar, V.; Zhang, H. Improving Fairness, Efficiency, and Stability in HTTP-Based Adaptive Video Streaming With Festive. *IEEE/ACM Trans. Netw.* **2014**, *22*, 326–340. [\[CrossRef\]](#)
8. Li, Z.; Zhu, X.; Gahm, J.; Pan, R.; Hu, H.; Begen, A.C.; Oran, D. Probe and Adapt: Rate Adaptation for HTTP Video Streaming At Scale. *IEEE J. Sel. Areas Commun.* **2014**, *32*, 719–733. [\[CrossRef\]](#)
9. Liu, J.; Tao, X.; Lu, J. QoE-Oriented Rate Adaptation for DASH With Enhanced Deep Q-Learning. *IEEE Access* **2019**, *7*, 8454–8469. [\[CrossRef\]](#)
10. Tran, C.M.; Nguyen Duc, T.; Tan, P.X.; Kamioka, E. FAURAS: A Proxy-Based Framework for Ensuring the Fairness of Adaptive Video Streaming over HTTP/2 Server Push. *Appl. Sci.* **2020**, *10*, 2485. [\[CrossRef\]](#)
11. *Hypertext Transfer Protocol Version 3 (HTTP/3)—Draft-Ietf-Quic-Http-34*; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2021.
12. *RFC: 7540 Hypertext Transfer Protocol Version 2 (HTTP/2)*; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2015.
13. *QUIC: A UDP-Based Multiplexed and Secure Transport—Draft-Ietf-Quic-Transport-34*; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2021.
14. Berti-Equille, L.; Zhauniarovich, Y. Profiling DRDoS Attacks with Data Analytics Pipeline. In Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17, Singapore, 6–10 November 2017; Association for Computing Machinery: New York, NY, USA, 2017; pp. 1983–1986. [\[CrossRef\]](#)

15. Wangen, G.; Shalaginov, A.; Hallstensen, C. Cyber Security Risk Assessment of a DDoS Attack. In *Information Security*; Bishop, M., Nascimento, A.C.A., Eds.; Springer International Publishing: Cham, Switzerland, 2016; pp. 183–202.
16. Soni, M.; Rajput, B.S. Security and Performance Evaluations of QUIC Protocol. In *Data Science and Intelligent Applications*; Kotecha, K., Piuri, V., Shah, H.N., Patel, R., Eds.; Springer: Singapore, 2021; pp. 457–462.
17. HTTP Archive's Annual State of the Web Report. Available online: <https://almanac.httparchive.org/en/2020/http2#fig-4> (accessed on 27 April 2021).
18. *QUIC Loss Detection and Congestion Control—Draft-Ietf-Quic-Transport-34*; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2021.
19. Arisu, S.; Begen, A.C. Quickly Starting Media Streams Using QUIC. In Proceedings of the 23rd Packet Video Workshop, PV '18, Amsterdam, The Netherlands, 12 June 2018; Association for Computing Machinery: New York, NY, USA, 2018; pp. 1–6. [CrossRef]
20. HTTP-over-QUIC Officially Becomes HTTP/3. Available online: <https://daniel.haxx.se/blog/2018/11/11/http-3/> (accessed on 9 March 2021).
21. Nguyen, M.; Amirpour, H.; Timmerer, C.; Hellwagner, H. Scalable High Efficiency Video Coding Based HTTP Adaptive Streaming over QUIC. In Proceedings of the Workshop on the Evolution, Performance, and Interoperability of QUIC, EPIQ '20, Virtual Event, USA, 10–14 August 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 28–34. [CrossRef]
22. Tong, V.; TRAN, H.A.; SOUHI, S.; MELLOUK, A. Empirical study for Dynamic Adaptive Video Streaming Service based on Google Transport QUIC protocol. In Proceedings of the 2018 IEEE 43rd Conference on Local Computer Networks (LCN), Chicago, IL, USA, 1–4 October 2018, pp. 343–350. [CrossRef]
23. Yen, S.C.; Fan, C.L.; Hsu, C.H. Streaming 360° Videos to Head-Mounted Virtual Reality Using DASH over QUIC Transport Protocol. In Proceedings of the 24th ACM Workshop on Packet Video, PV '19, Amsterdam, The Netherlands, 12–15 June 2018; Association for Computing Machinery: New York, NY, USA, 2019; pp. 7–12. [CrossRef]
24. Bhat, D.; Rizk, A.; Zink, M. Not so QUIC: A Performance Study of DASH over QUIC. In Proceedings of the 27th Workshop on Network and Operating Systems Support for Digital Audio and Video, Taipei, Taiwan, 20–23 June 2017; pp. 13–18. [CrossRef]
25. Bhat, D.; Deshmukh, R.; Zink, M. Improving QoE of ABR Streaming Sessions through QUIC Retransmissions. In Proceedings of the 26th ACM International Conference on Multimedia, Seoul, Korea, 22 October 2018. [CrossRef]
26. Arisu, S.; Yildiz, E.; Begen, A.C. Game of Protocols: Is QUIC Ready for Prime Time Streaming? *Int. J. Netw. Manag.* **2020**, *30*. [CrossRef]
27. Chrome Is Deploying HTTP/3 and IETF QUIC. Available online: <https://blog.chromium.org/2020/10/chrome-is-deploying-http3-and-ietf-quic.html> (accessed on 27 April 2021).
28. *RFC: 5681 TCP Congestion Control*; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2009.
29. Ha, S.; Rhee, I.; Xu, L. CUBIC: A New TCP-Friendly High-Speed TCP Variant. *Acm Sigops Oper. Syst. Rev.* **2008**, *42*, 64–74. [CrossRef]
30. Jain, R.; Chiu, D.; Hawe, W. A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems. *arXiv* **1998**, arXiv:cs/9809099.
31. quic-go. Available online: <https://github.com/lucas-clemente/quic-go/tree/master> (accessed on 9 March 2021).
32. golang http. Available online: <https://golang.org/pkg/net/http/> (accessed on 9 March 2021).
33. linux tc. Available online: <https://man7.org/linux/man-pages/man8/tc.8.html> (accessed on 9 March 2021).
34. linux ss. Available online: <https://man7.org/linux/man-pages/man8/ss.8.html> (accessed on 9 March 2021).
35. dash.js. Available online: <https://github.com/Dash-Industry-Forum/dash.js/wiki> (accessed on 9 March 2021).
36. Spiteri, K.; Urganikar, R.; Sitaraman, R.K. BOLA: Near-optimal bitrate adaptation for online videos. In Proceedings of the IEEE INFOCOM 2016, 35th Annual IEEE International Conference on Computer Communications, San Francisco, CA, USA, 10–14 April 2016; pp. 1–9.
37. Taibi Guguen, C.; Le Bolzer, F.; Houdaille, R. Improving User Experience when HTTP Adaptive Streaming Clients Compete for Bandwidth. *Smpete Motion Imag. J.* **2017**, *126*, 28–34. [CrossRef]
38. Bentaleb, A.; Taani, B.; Begen, A.C.; Timmerer, C.; Zimmermann, R. A Survey on Bitrate Adaptation Schemes for Streaming Media Over HTTP. *IEEE Commun. Surv. Tutor.* **2019**, *21*, 562–585. [CrossRef]
39. Petrangeli, S.; Hooft, J.V.D.; Wauters, T.; Turck, F.D. Quality of Experience-Centric Management of Adaptive Video Streaming Services: Status and Challenges. *ACM Trans. Multimed. Comput. Commun. Appl.* **2018**, *14*. [CrossRef]
40. Wang, P.; Bianco, C.; Riihijärvi, J.; Petrova, M. Implementation and Performance Evaluation of the QUIC Protocol in Linux Kernel. In Proceedings of the 21st ACM International Conference on Modelling, Analysis and Simulation of Wireless and Mobile Systems, MSWIM '18, Montreal, QC, Canada, 28 October–2 November 2018; Association for Computing Machinery: New York, NY, USA, 2018; pp. 227–234. [CrossRef]
41. Kosek, M.; Shreedhar, T.; Bajpai, V. Beyond QUIC v1: A First Look at Recent Transport Layer IETF Standardization Efforts. *IEEE Commun. Mag.* **2021**, *59*, 24–29. [CrossRef]
42. Papastergiou, G.; Fairhurst, G.; Ros, D.; Brunstrom, A.; Grinnemo, K.J.; Hurtig, P.; Khademi, N.; Tüxen, M.; Welzl, M.; Damjanovic, D.; et al. De-Ossifying the Internet Transport Layer: A Survey and Future Perspectives. *IEEE Commun. Surv. Tutor.* **2017**, *19*, 619–639. [CrossRef]