Temporal Characterization of XR Traffic with Application to Predictive Network Slicing

Mattia Lecci*, Federico Chiariotti†, Matteo Drago*, Andrea Zanella*, and Michele Zorzi*

*Department of Information Engineering, University of Padova, 35131 Padua, Italy

†Department of Electronic Systems, Aalborg University, 9220 Aalborg Øst, Denmark

Emails: *{leccimat, dragomat, zorzi}@dei.unipd.it, †fchi@es.aau.dk

Abstract—Over the past few years, eXtended Reality (XR) has attracted increasing interest thanks to its extensive industrial and commercial applications, and its popularity is expected to rise exponentially over the next decade. However, the stringent Quality of Service (QoS) constraints imposed by XR's interactive nature require Network Slicing (NS) solutions to support its use over wireless connections: in this context, quasi-Constant Bit Rate (CBR) encoding is a promising solution, as it can increase the predictability of the stream, making the network resource allocation easier. However, traffic characterization of XR streams is still a largely unexplored subject, particularly with this encoding. In this work, we characterize XR streams from more than 4 hours of traces captured in a real setup, analyzing their temporal correlation and proposing two prediction models for future frame size. Our results show that even the state-of-the-art H.264 CBR mode can have significant frame size fluctuations, which can impact the NS optimization. Our proposed prediction models can be applied to different traces, and even to different contents, achieving very similar performance. We also show the trade-off between network resource efficiency and XR QoS in a simple NS use case.

Index Terms—Virtual Reality, Extended Reality, Traffic Modeling, Network Slicing, Resource Provisioning

I. INTRODUCTION

Over the past few years, the rapid technological development of Head Mounted Devices (HMDs) and the strong push towards the virtual world caused by the COVID-19 pandemic have caused an explosion of the eXtended Reality (XR) market, which includes technologies such as Virtual Reality (VR), Augmented Reality (AR), and Mixed Reality (MR). Recent studies estimate hundreds of millions of users of these technologies in a time span of just 3 years [1], requiring millions of new devices to be developed, produced, and shipped around the world for a business in the order of billions of dollars [2].

While the latest news on the *metaverse* seem to indicate that the fastest growth will be in the entertainment and social media industries, XR is expected to make an impact in a huge variety of scenarios [3], [4]. Interactive design, marketing, healthcare, and employee training are just a few of the proposed scenarios, but industrial remote control in manufacturing and agriculture might have the largest impact, allowing human operators to remotely control machines in risky, hard to reach or unsafe environments, through a fully interactive virtual framework.

One of the peculiarities shared by all these new applications is their interactive nature: users do not passively receive the information or stream a video, but need to manipulate the environment and affect it in meaningful ways, while maintaining an illusion of presence which requires the application to operate under very strict end-to-end delay constraints [5], [6]. In particular, safety-critical and industrial applications will have stricter constraints, as the consequences of network impairments can be significantly more serious. *Cybersickness* is another important issue, as a delay over 20 ms between movements and visual and auditory feedback can cause disorientation and dizziness [1], [7].

In order to fulfill these stringent latency requirements over a wireless connection, the application and the network need to cooperate. The *Network Slicing (NS)* paradigm [8] allows 5G and Beyond networks to reserve resources to a given stream, defining Quality of Service (QoS) targets, but most efforts in this sense focus on relatively predictable applications. In this setting, the need for predictability in the XR traffic becomes extremely important, leading to a resurgence of quasi-Constant Bit Rate (CBR) encoders, which are not used in passive streaming due to their lower picture quality stability. While some efforts have been devoted by prominent standard bodies on this topic [5], [6], the current availability of traffic models for XR is scarce. Furthermore, to the best of our knowledge, no detailed analysis of the temporal statistics of quasi-CBR video streams can be found in the literature, making existing scheduling schemes rely on uncertain foundations.

However, even CBR encoders are not perfect, and the interplay between the video content and the movements and actions of the users may cause significant fluctuations. In this work, we analyze the traffic from a real VR application using the Periodic-Intra Refresh mode of the H.264 codec, which results in relatively small differences in the frame sizes. Modeling these imperfections, and consequently predicting the size of future frames in advance, can be extremely significant in the allocation of network resources, particularly if some critical QoS metrics have to be reached. For example, this is the case for *Cloud XR*, a new trend pursued by some major players in the telecommunication industry that moves the processing and rendering steps of the XR content from the user to the Cloud, making the QoS requirements even more

critical [9], [10].

In this paper, we hence address the problem of providing a realistic stochastic characterization of an XR traffic source. Building upon our previous works [11], [12] where we collected more than 4 hours of live sessions with a real HMD and performed basic traffic characterization, in this paper we take the analysis one step further by modeling the size of XR frames in the stream as a correlated time series. We propose two parametric regression models to predict the size of future frames, and show that the behavior of the encoder can be generalized to other traces and even different applications with limited regression performance loss. Finally, we present a simple network slicing use case, in which we show the trade-off between resource efficiency and latency for different types of resource scheduling. All our traces as well as the analysis and simulation code is publicly available.[1]

The rest of the paper is structured as follows. Sec. II will discuss the current state of the art on XR modeling, and our experimental setup and VR application are briefly presented in Sec. III. Our analysis is reported in Sec. IV, while Sec. V illustrates how our analysis can be leveraged for a simple NS use case. Finally, Sec. VI draws conclusions and presents some avenues for future work.

II. STATE OF THE ART

Despite a steady scientific interest in VR since the 1990s [13], relatively little work has been done to characterize the details of this type of traffic. With respect to our work, we can distinguish two main areas of research: the modeling and characterization of XR traffic, and the scheduling and resource management of XR data streams.

The former is closely related to 2D video content, and, even more so, to live, interactive applications such as video conferencing and gaming. However, most of the work on the subject has focused on Variable Bit Rate (VBR) encoding, based either on the H.264 or the H.265 standards [14], i.e., the customary encoding for streaming pre-generated video content. VBR can provide a stable visual quality, improving the user Quality of Experience (QoE), but is also subject to significant jitter due to the large frame size fluctuations. Transmitting VBR videos with low latency can then be a significant challenge even over channels with constant capacity [15]. On the other hand, CBR encoding sacrifices some visual quality stability to obtain an encoded video stream with a stable transmission rate [16]. Although the higher predictability of the encoded output makes CBR encoding attractive for interactive video and XR content, it is still relatively unexplored in the relevant literature.

A topic related to XR traffic is video game streaming, also called *Cloud gaming*. These Cloud frameworks run games over a remote server, streaming the screen directly to the users without the need for client-side computation. The stringent requirements of gaming applications, especially in terms of latency, and the need to address them with optimized protocols

and new transmission strategies, have led to an increased interest in their characterization.

The authors of [17] carried out an extensive measurement campaign in Google Stadia, a famous cloud gaming platform, giving an overview of its inner workings. They studied the distributions of downlink traffic, packet size, inter-packet time under multiple settings, including different resolutions, video codecs, and network conditions. On the other hand, in [18], [19] direct comparisons were made between different cloud gaming platforms, mostly focusing only on the bitrate of the video stream, without including latencies or user experience.

A more comprehensive Cloud gaming testbed, including a full implementation with different network alternatives and automated trace acquisition over Ethernet, WiFi, and LTE, was presented in [20]. This is surely an advantage in terms of reproducibility and speed of the experiments, but the unpredictability of the users' actions in gaming scenarios (and, more importantly, in XR) is the real challenge that the network has to face, limiting the usefulness of the results.

These works represent a good starting point for the collection and modeling of XR traffic, as it is reasonable to assume that most of these Cloud gaming companies will start providing XR services soon. However, most works focused specifically on XR still consider simple applications, such as interactive data visualization [21], and do not provide much insight on the more complex scenarios. There is an extensive literature on immersive video streaming [22], but it has been mostly focused on passive applications in which the user is only a viewer, with different QoE and encoding considerations.

Regarding XR traffic scheduling and resource management, some works have already tried to propose schemes for efficient systems. For example, in [23], [24] game-theoretic approaches are proposed to tackle the optimization of multi-user VR streaming over a small cell, with the help of machine learning. The authors of [25] analyze the scheduling problem from the perspective of Mobile Edge Cloud (MEC), proposing scheduling strategies and analyzing communication, computing, and caching trade-offs. While the models proposed for the network architectures considered in these works are extremely complex, there is no comparison with real-world VR streaming.

To the best of our knowledge, our previous works, which proposed a simple architecture for collecting traffic traces from VR games [11] and a simple generative model for the frame size [12], were the first to use real VR traffic traces. This work extends our previous ones by characterizing the temporal behavior of the XR traces and drawing novel conclusions for NS optimization.

III. XR STREAMING ARCHITECTURE

In this section, we describe the architecture of our XR streaming acquisition and give some perspective on the full end-to-end setup. To further understand what are the steps that most influence the XR performance, it is useful to describe a common end-to-end XR architecture. First, we can start from the collection and processing of tracking information, delegated to the HMD. Then, this information is sent to a remote

---

[1]Code repository: https://github.com/signetlabdei/vr-trace-analysis
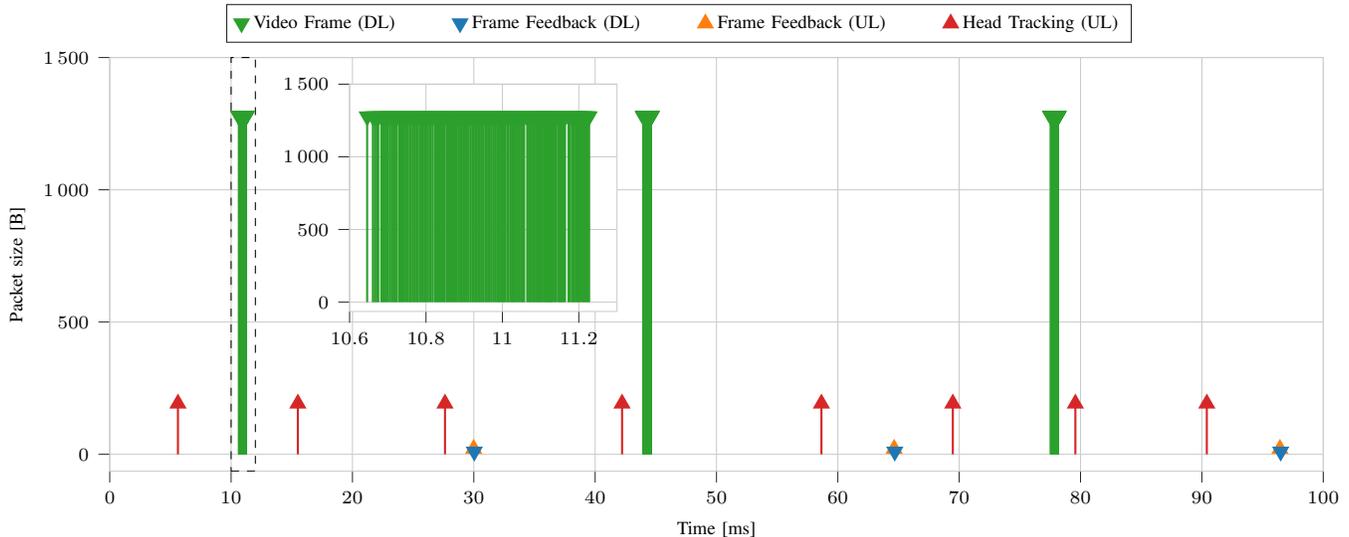
Fig. 1: Portion of traffic trace from *Virus Popper* (50 Mb/s, 30 FPS). For this trace, about 130–140 individual fragments make up each video frame burst.

server to compose the viewport, i.e., what is actually shown to the user. This process includes the rendering of the scene, the video encoding providing a more robust transmission towards the mobile device, and possibly some additional information e.g., the direction in which the rendered frame is supposed to be displayed. After receiving and decoding the video stream together with all the additional meta-information, the HMD generates the images to display at the occurring screen refresh. These steps need to be accomplished with minimal delay to guarantee adequate QoE.

Our experimental setup consisted of a desktop computer equipped with an NVIDIA GeForce RTX 2080 Ti graphics card acting as the rendering server, and an iPhone XS enclosed in a VR cardboard acting as the HMD. VR applications were thus run on the rendering server and streamed to the headset using the *RiftCat 2.0* application (on the server), and *VRidge* 2.7.7 (on the phone).[2]

The application uses hardware-accelerated H.264 encoding via Nvidia Encoder (NVENC) as long as a compatible graphics card is present on the system. *RiftCat*'s developers disclosed that Periodic Intra-Refresh is used, a setting provided by the encoder that allows each frame to be roughly the same size, making the stream almost CBR and thus easier to handle from a network perspective. It does so by replacing key-frames by *waves* of refreshed intra-coded blocks, i.e., blocks without any dependence on other frames, effectively spreading a key frame over multiple frames. Image quality is balanced with resilience to packet loss by setting the `intraRefreshPeriod` parameter, which determines the period after which an intra refresh happens again, and the `intraRefreshCnt` parameter, which sets the number of frames over which the intra refresh happens [26]. If we consider a 30 Frames per Second (FPS) video, a value of 30 for the `intraRefreshPeriod` would

ensure that the frame is fully recovered every second. On the other hand, choosing the value of `intraRefreshCnt` determines the number of frames over which the intra refresh will happen within an intra refresh period, with smaller values leading to a quicker refresh but lower quality.

Detailed information about the video encoder is of the utmost importance for our work, since different encoders typically behave differently, especially when analyzing the temporal behavior of the encoded source. Still, we believe that our work offers network researchers a peek into the intricacies of this topic, showing some key results on how an XR traffic flow can be analyzed for resource provisioning.

Different freely available games and applications were used to acquire our dataset, including *Minecraft*, *Virus Popper*, and *Google Earth VR*. Further details on the acquisition setup and our traces can be found in [12]. In the following, we will mostly concentrate on one trace acquired using the *Virus Popper* application, but the methodology holds throughout the dataset, and can be easily replicated for any of the other traces.

IV. VIDEO TRACE ANALYSIS

Analyzing the acquired traces, we determined that the application used User Datagram Protocol (UDP) over IPv4. It also used an additional application-layer protocol header of variable size, which we decoded to determine the types of packets being exchanged. More specifically, synchronization and acknowledgment packets were exchanged in both directions, while the Uplink (UL) stream from the HMD to the rendering server also contained frequent and relatively small head-tracking information packets. Naturally, the Downlink (DL) stream also had regular video frame packet bursts.

Fig. 1 is a visual representation of a slice of bidirectional VR streaming, showing the main data streams in both DL and UL. As the figure clearly shows, most of the traffic is

---

(a) Rate distribution with different moving average windows.

(b) Overflow rate for different moving average windows.
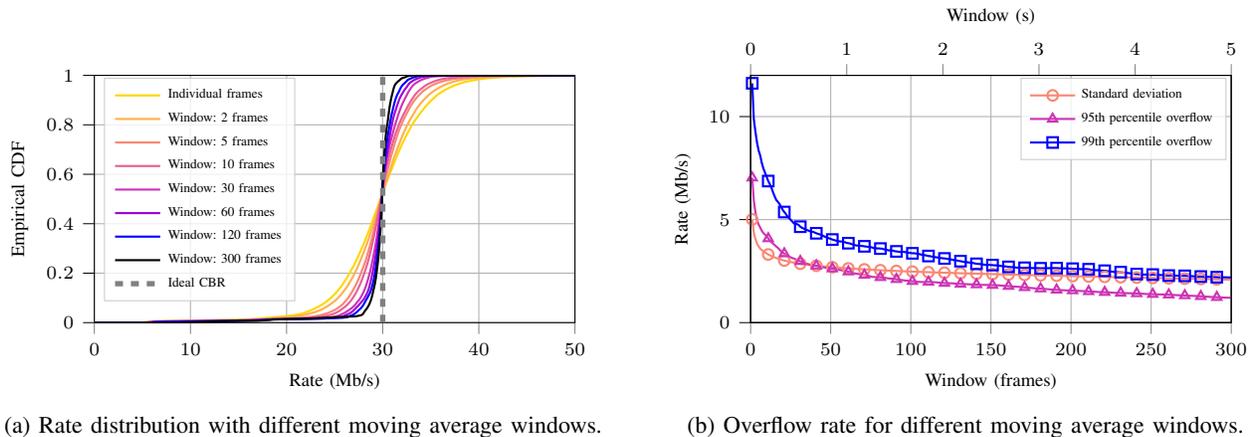
Fig. 2: Video frame size statistics for *Virus Popper* (30 Mb/s, 60 FPS).

concentrated in DL and is made up of packet bursts encoding video frames. Video frame fragments were consistently found to be 1320 B long in all acquired traces, with a data size (the UDP payload) of 1278 B.

The low impact of non-video packets on the total streaming data rate, along with their strong dependence on the application setup, led us to focus exclusively on the video frame data, discarding all other packets from our analysis. Our results can then be applied to any VR application using the same encoder. By decoding the application protocol, we managed to identify frame boundaries and extract the video frame data, removing metadata and control information. We can then consider the size of individual frames in a video trace.

The encoder makes use of the H.264 Periodic Intra-Refresh compression scheme to reduce the variation between frame sizes, so we do not expect a multimodal distribution, as would be the case for a classical keyframe-based encoding. As we mentioned above, encoding VR traffic as CBR can be significantly better for network optimization, although it leads to a less stable picture quality: if all frames have the same size, it is possible for network slicing schemes to provide a guaranteed latency without wasting resources.

However, CBR encoding is not perfect, and frames may still have variable size, although the average rate almost perfectly matches the required one. We can use a simple Moving Average (MA) filter to examine the behavior of the VR traffic on longer timescales, which is useful if resource allocation is performed at a slower pace. Naturally, allocating resources every N frames leads to a larger jitter between frames, but it can also improve the resource allocation efficiency, as size fluctuations will tend to average out over multiple frames.

In order to measure this effect, we consider the *Virus Popper* trace, with a required rate $R = 30$ Mb/s and a $\varphi = 60$ FPS refresh rate. We only measure the video traffic, without packet headers and redundancy added by the application: this results in an average rate of 29.76 Mb/s. Fig. 2a shows the empirical Cumulative Distribution Function (CDF) of the rate, considering different MA window sizes. If we consider each frame
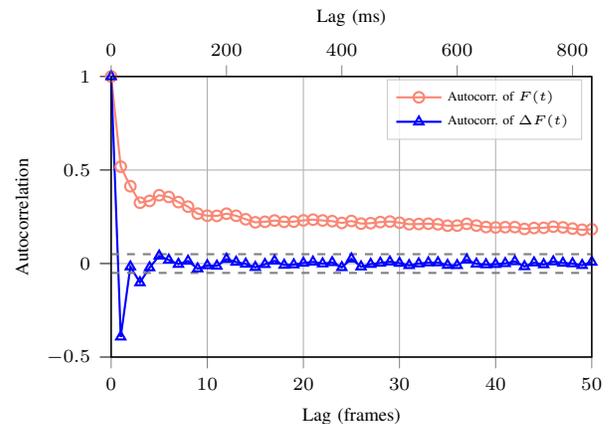
Fig. 3: Video frame size autocorrelation for *Virus Popper* (30 Mb/s, 60 FPS).

individually, there is a significant variation, which gradually reduces as we increase the period over which the rate is measured. It is also possible to notice that the frame size distribution is skewed towards smaller frame sizes, even for longer windows, as can be seen from the asymmetric tails of the distributions.

However, providing reliable service will require a significant overhead even if we relax the scheduling time: Fig. 2b shows the overflow rate (i.e., the difference between the actual rate and the expected 30 Mb/s CBR rate) as a function of the MA window. The plot shows the standard deviation, as well as the 95[th] and 99[th] percentile overflow rates. If our aim is to provide 99% reliability, we need to overprovision by more than 8 Mb/s (i.e., almost 30% of the CBR rate) even if we consider a timescale of 100 ms for resource allocation, i.e., 6 frames. Even averaging over periods of multiple seconds leads to worst-case rates almost 4 Mb/s higher than the average, probably corresponding to highly dynamic content in the video or to how the CBR encoder works. Interestingly, the standard deviation does not decay as fast as the 95[th] and
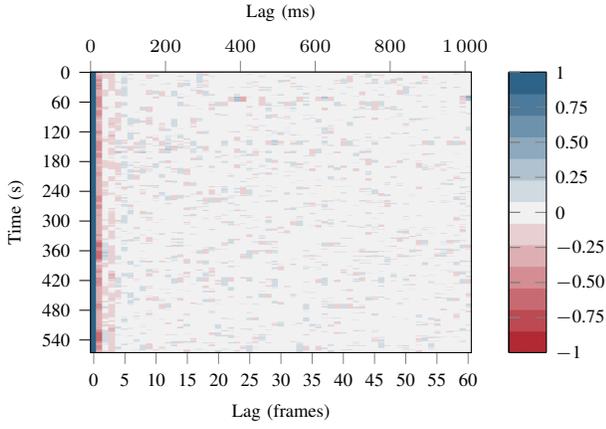
Fig. 4: Rolling windowed ΔF autocorrelation for *Virus Popper* (30 Mb/s, 60 FPS). The windows were 600 frames (10 s) long, with a time shift of 60 frames (1 s).

99^{th} percentile overflows for longer averaging windows, due to the fact that the frame size distribution is skewed towards smaller sizes, as previously highlighted.

We can also analyze the autocorrelation of the frame size signal $F(t)$, to identify patterns in how the signal changes. Fig. 3 shows the autocorrelation of $F(t)$ and $\Delta F(t) = F(t) - F(t-1)$. While $F(t)$ has a strong long-term autocorrelation, due to the constant component, the $\Delta F(t)$ signal has a strong negative autocorrelation between one frame and the next, while almost all longer time differences fall within the ± 0.05 range.

This means that the encoder tends to balance out fluctuations between one frame and the next, such that a frame that is bigger than the previous one tends to be followed by a smaller one again. We can check that this holds throughout the whole video by computing a rolling window autocorrelation, showed in Fig. 4 for $\Delta F(t)$. In this case, the plot clearly shows that there are no strong long-term correlations in any part of the video. The frame difference signal has a noticeable autocorrelation only with lag 1 and 3, confirming the result from Fig. 3.

A. Frame Size Prediction

Let us consider the average size of future frames in the time interval $[t, t+T)$, given by

$$\overline{F_T}(t) = \frac{1}{T} \sum_{i=0}^{T-1} F(t+i). \qquad (1)$$

We denote by $\hat{F}_T(t, \tau)$ an estimate of $\overline{F_T}(t+\tau)$, $\tau > 0$, i.e., considering a look-ahead of τ frames. We focus on linear predictors based on the last $N \geq 0$ samples, so that

$$\hat{F}_T(t, \tau) = \theta_0 + \sum_{j=1}^{N} \theta_j F(t-j+1), \qquad (2)$$

where $\boldsymbol{\theta} = [\theta_0, \ldots, \theta_N]$ is a weight vector, which determines the accuracy of the estimate. The difference between actual and estimated value is captured by the error term $w(t, \tau, T) =$ $\overline{F_T}(t+\tau) - \hat{F}_T(t, \tau)$, which will be denoted just as w in the following, for ease of writing.

We can then consider two different regression methods to determine the value of the parameter vector $\boldsymbol{\theta}$:

- *Ordinary Least Squares (OLS) linear regression*: least squares regression was independently developed by Gauss and Legendre in the 19$^{\text{th}}$ century [27], and is the most classic form of regression. In this case, the objective is to minimize the ℓ^2 norm of the signal w. OLS regression can be useful in determining the average behavior of the underlying stochastic process, giving easily interpretable results on the quality of the prediction and the dynamics of the frame size over time;
- *Quantile regression* [28]: this technique estimates $\hat{F}_T(t, \tau)$ so that the probability that it is higher than the real value, is not larger than p_s. This has obvious implications for the main objective of this paper, which is VR traffic modeling for network resource provisioning: as we are interested in providing enough resources to send a frame within the required latency with probability p_s, estimating the corresponding quantile might be the best way to get the required quality.

We also used *Robust linear regression* [29] to verify that the OLS prediction was not too sensitive to outliers. We considered a robust method using Huber's T norm instead of the ℓ^2 norm: the two norms have the same quadratic behavior if the error is smaller than a threshold δ, but Huber's T increases linearly for larger values. Setting the threshold to $\delta = \frac{\mathbb{E}[|F|]}{4}$, we found that the results matched exactly those of the OLS model, suggesting that outliers are not playing a relevant role in this case and thus letting us discard this model.

In this section, we will show results for both the OLS and the quantile regression models. As we stated above, while the results from OLS are more immediate, quantile regression is useful when focusing on scheduling network resources for a VR stream, which requires a model of the tail of the frame size distribution to provide latency guarantees.

We can now examine the results of the regression analysis for the *Virus Popper* trace, considering a rate of 30 Mb/s and 60 FPS. We focus on this video trace as the standard example in the paper, but other traces, even at different bitrates and frame rates, exhibit a similar behavior. Fig. 5 shows the complementary CDF of the residual error w, considering $\tau = 1$ and two different values of T. The first thing we can notice is that the error distribution has a slightly different shape for the OLS and quantile regression models, indicating that the difference in the two models is not simply a shift in the value of the intercept θ_0, but instead the two predictions are meaningfully different. We can also notice that there is some benefit from having a longer memory, although increasing N yields diminishing returns. Finally, we can confirm that the reliable transmission of this VR content will require significant overprovisioning, even when using prediction: for $T = 1$ the 95th percentile error of the OLS prediction is approximately 15 kB higher than the mean with any of the models, i.e., about 25% of the average frame size (which is 62.5 kB for

(a) OLS regression, $T = 1$.

(b) Quantile regression ($p_s = 0.95$), $T = 1$.

(c) OLS regression, $T = 6$.

(d) Quantile regression ($p_s = 0.95$), $T = 6$.
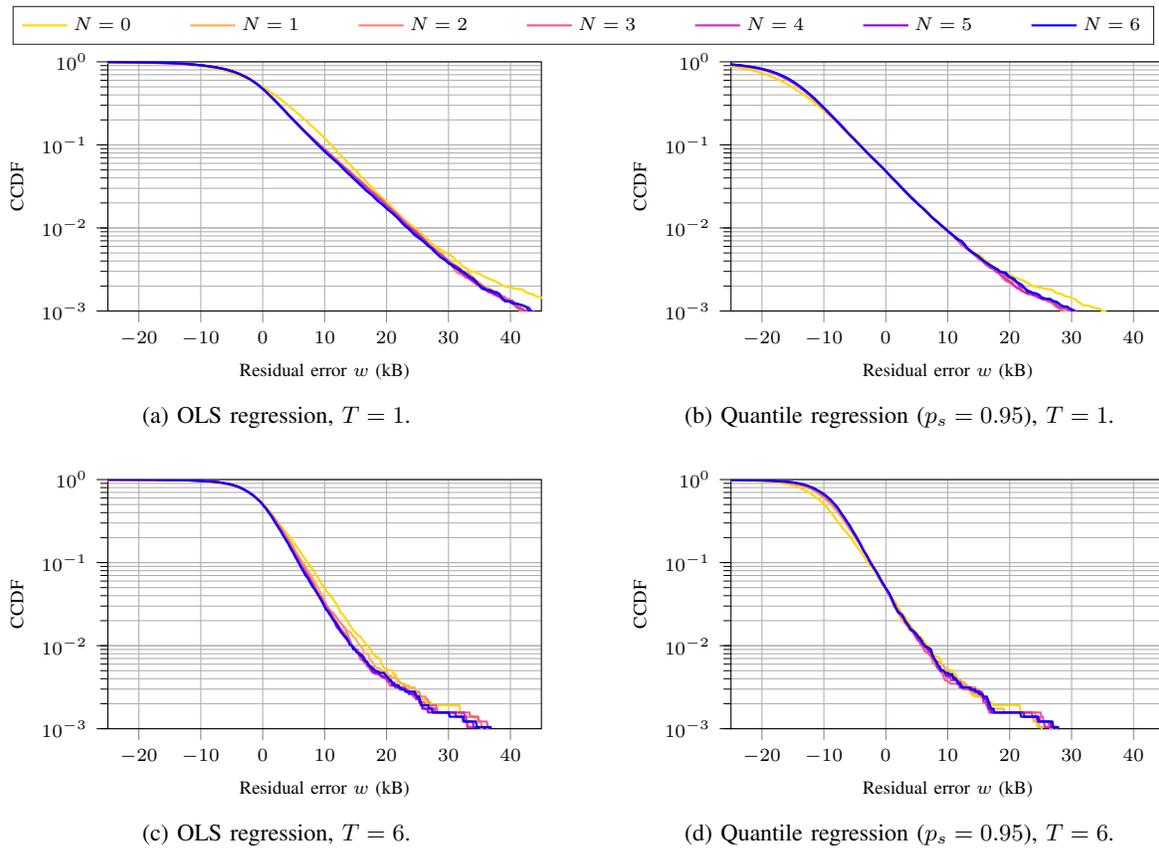
Fig. 5: Complementary CDF of the residual error w with $\tau = 1$, for different values of N and T.

(a) OLS regression.

(b) Quantile regression ($p_s = 0.95$).
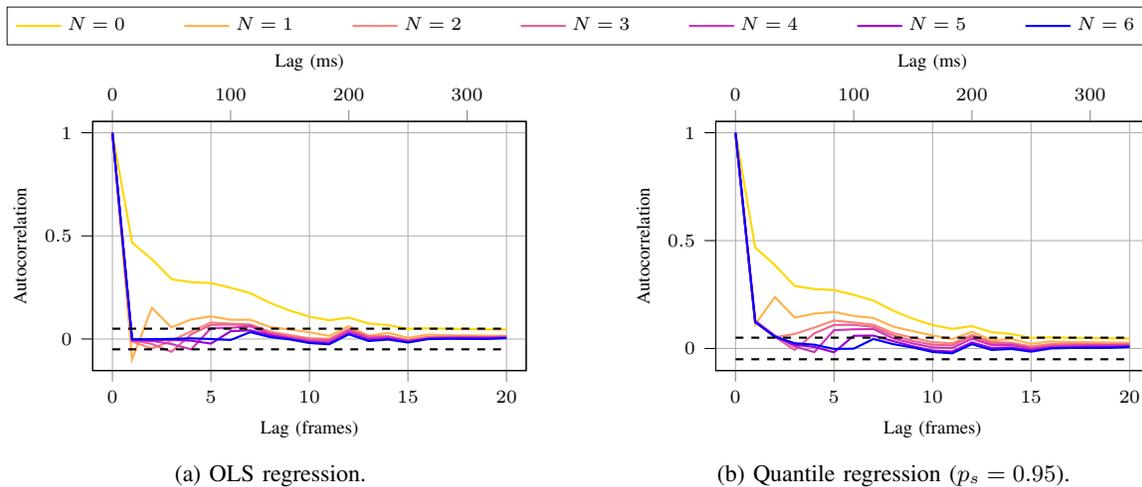
Fig. 6: Autocorrelation of the residual error w for next-frame prediction ($T = 1$, $\tau = 1$) for different values of N.

this trace). In fact, this is close to the difference between the average predictions of the OLS and quantile models.

This difference is about halved for $T = 6$, due to the fact that computing the average over multiple frames allows errors to compensate and cancel each other out. However, provisioning over multiple frames means that only the average amount of resources will be scheduled for the stream, which will cause larger frames to have a higher latency, thus taking

longer than $\frac{1}{\varphi}$ seconds to be delivered and causing additional queuing delay to subsequent frames. Since the frame cannot be properly shown on screen until it is fully received, this translates to a higher jitter and reduces the QoE perceived by the user, making a lower value of T preferable.

Another fundamental component in evaluating the quality of a predictor is the autocorrelation of the residual error w: if the autocorrelation between subsequent samples of the residual

(a) OLS regression. (b) Quantile reg. ($p_s = 0.95$).
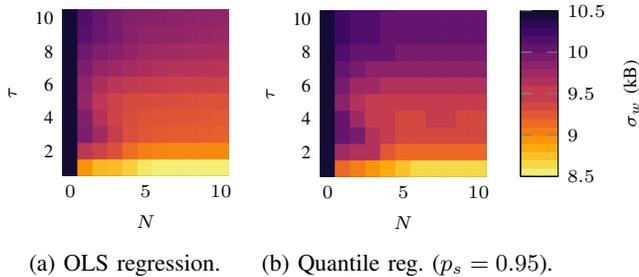
Fig. 7: Heatmap of the residual error standard deviation (measured in kB) as a function of N and τ, with $T = 1$.

error is high, the model did not capture some effect, usually due to an insufficient memory, i.e., too low a value of N. Fig. 6 shows the autocorrelation of w for different values of N: it is easy to see that models with $N < 4$, and particularly with $N = 0$ and $N = 1$, do not have enough memory to capture the frame size dynamics. This is more evident in quantile regression, which shows a higher autocorrelation for these models.

Finally, we can examine the effect of N and τ on the quality of the prediction by looking at Fig. 7, which shows the standard deviation of the residual error w as a function of these two parameters with $T = 1$. The figure clearly shows that increasing the memory of the model improves the prediction, but gives diminishing returns, as the difference between $N = 6$ and $N = 10$ is minimal. Furthermore, we see an expected increase in the error if τ increases, but this is not monotonic for $N < 3$: this might be due to the autocorrelation we observed in the w signal, as $N < 3$ is not sufficient to fully represent the state of the stochastic process, resulting in suboptimal predictions.

B. Model Generalization

In the above, we studied how well regression models can predict future frame sizes $\hat{F}_T(t, \tau)$, but we always found the parameter vector $\boldsymbol{\theta}$ based on the same video trace. In the following, we study how prediction models perform when the regression is performed over multiple traces, with different bitrates and types of content. This has significant advantages, as finding a predictor for each specific video content requires acquiring traces for each content and quality level, while generalizing the predictor would allow for simpler deployment.

We consider $N = 6$ and $\tau = 1$, as we determined that $N = 6$ is sufficient to capture the dynamics of the model. In order to directly compare traces with different bitrates R and frame rates φ, we normalize the video traces by the expected frame size $\varphi^{-1}R$, obtaining a normalized parameter vector $\tilde{\boldsymbol{\theta}}$, which, given the linearity of our models, can be converted back to the original parameter vector as $\boldsymbol{\theta} = \frac{R\tilde{\boldsymbol{\theta}}}{\varphi}$ of the regression model in (1). By normalizing our frame sizes, we can train and use our models on multiple traces with different values of R and φ. We then consider three generalized models:

1) A *general* model (GM), which computes $\boldsymbol{\theta}$ using the whole dataset, with different frame rates, bitrates, and video content types;

2) A *content-dependent* model (CM), which computes $\boldsymbol{\theta}$ using a single type of content (e.g., the *Virus Popper* game), but with different bitrates and frame rates;
3) A *content- and rate-dependent* model (CRM), which derives the parameter vector on a per-content, frame rate, and bitrate basis, i.e., a single trace.

Given that different values of R and φ can have different scales of errors which can be difficult to compare directly, in Fig. 8 we show the error normalized to the expected frame size R/φ. As the figure shows, the model can generalize quite well: the performance of CM is almost always similar to that obtained by CRM, making generalization across different bitrates and frame rates possible for the same video content. On the other hand, GM performs slightly worse, and has a large error in the Minecraft trace with $R = 40$ Mb/s: it is possible that this trace involves different dynamics in the content or head movements, leading to sharp differences even with other traces with the same type of content. On the other hand, GM has similar performance to CM and CRM with the OLS predictor, but shows a less consistent behavior for the quantile regressor. For example, the *Minecraft* trace with $R = 40$ Mb/s shows very different performance between the three models and different values of T. Furthermore, the *Virus Popper* trace seems to have a smaller tail, as GM is more conservative than the models based only on that video content.

As we can see, using the quantile model leads to the prediction being between 25% and 40% higher than the average, skewing the error distribution. We can also note that the relative error decreases with the bitrate: lower bitrate traces have a higher prediction error relative to the frame size, although the raw error w is still larger for increasing bitrates. As we noted above, averaging over multiple frames can also significantly reduce the error across almost all traces.

V. NETWORK SLICING USE CASE

In this section, we consider an NS use case for the models we developed in Sec. IV. The VR stream is assigned to a high-priority slice, with the objective of allowing each frame to be delivered before the generation of the next one, i.e., maintaining a latency below $1/60^{\text{th}}$ of a second. Provisioning the time and frequency resources for VR is a critical component of Beyond 5G networks, and guaranteeing limited latency while reducing the impact on other users is an important application of our model.

We can then assume that the network slicing orchestrator is equipped with the CM quantile regression from Sec. IV-B, and can predict the frame sizes for arbitrary values of T and τ. We consider an orchestrator that can make decisions on the resource allocation only at times $t = kS$, $k \in \mathbb{Z}$, i.e., every S frames or, conversely, every $\Delta t = \frac{S}{\varphi}$ ms. In the following, we consider queued bits from earlier frames in the scheduling as well. At time $t = kS$, we consider that the previous slice might have been unable to send all the data in time, leaving in the queue q_t bits that have to be sent in the following slices with an excess capacity of q_t/φ.

(a) OLS regression, $T = 1$.

(b) Quantile regression ($p_s = 0.95$), $T = 1$.

(c) OLS regression, $T = 6$.

(d) Quantile regression ($p_s = 0.95$), $T = 6$.
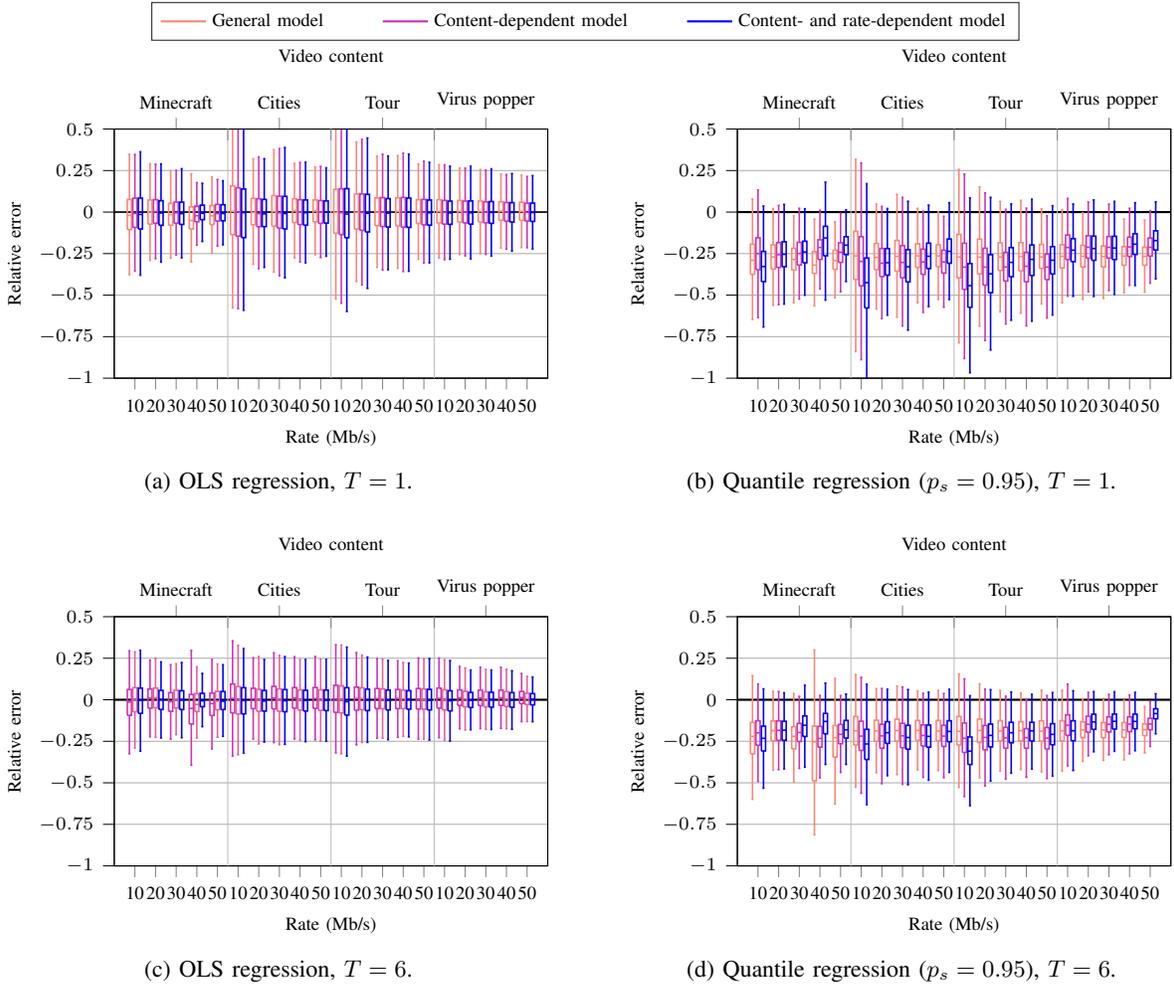
Fig. 8: Boxplot of the relative residual error $\frac{\varphi w}{R}$ for different levels of generalization with $N = 6$ and $\tau = 1$.

Recalling that our predictors are able to estimate values for $\hat{F}_T(t, \tau)$, as expressed in (1), we thus propose two different models:

1) *Constant* scheduling (CS), which only allows the scheduler to set a constant slice capacity $C(t)$ for the next S frames, i.e., only one prediction is performed, with $T = S$ and $\tau = 1$:

$$C_{\text{CS}}(kS + \ell) = \hat{F}_S(kS, 1) + \frac{q_t}{S\varphi}, \quad \ell = 1, \ldots, S, \quad (3)$$

where the excess capacity from the queued bits is spread among the following S frames;

2) *Frame-by-frame* scheduling (FS), in which a different slice capacity $C(t)$ can be set for every inter-frame period in the next S frames, i.e., there are S independent predictions, with $T = 1$ and $\tau \in \{1, \ldots, S\}$:

$$C_{\text{FS}}(kS + \ell) = \begin{cases} \hat{F}_1(kS, \ell) + \frac{q_t}{\varphi}, & \ell = 1; \\ \hat{F}_1(kS, \ell), & \ell = 2, \ldots, S, \end{cases} \quad (4)$$

where the excess capacity from the queued bits is added entirely to the next frame to minimize latency.

As we remarked in the previous section, the FS scheme can reduce the jitter by having a more fine-grained prediction, as each frame will be allocated enough resources to be transmitted with probability p_s. On the other hand, the CS scheme has a rougher prediction, with consequently higher jitter, but will waste fewer network resources, as it can allow larger frames to be compensated by smaller ones before and after them. Both models are realistic, as they work under different assumptions: in the first case, the resources that are allocated for each frame need to be over both time and frequency, while the second case gives the slice a constant bandwidth over the scheduling interval, the most common slicing model in the literature.

We can then look at the schedulers' performance as a function of S, setting $p_s = 0.95$ and $N = 6$: Fig. 9 shows boxplots of the latency and scheduled capacity for FS and CS. Fig. 9a clearly shows that, while the scheduler granularity has a limited effect on FS, the lower precision of CS means that the longer the scheduling interval, the higher the average latency, and the worst-case latency, represented by the upper whisker of the boxplots, increases even more. On the other hand, as Fig. 9b shows, the capacity required by CS decreases

(a) Quantile regression ($p_s = 0.95$), latency.

(b) Quantile regression ($p_s = 0.95$), scheduled capacity.
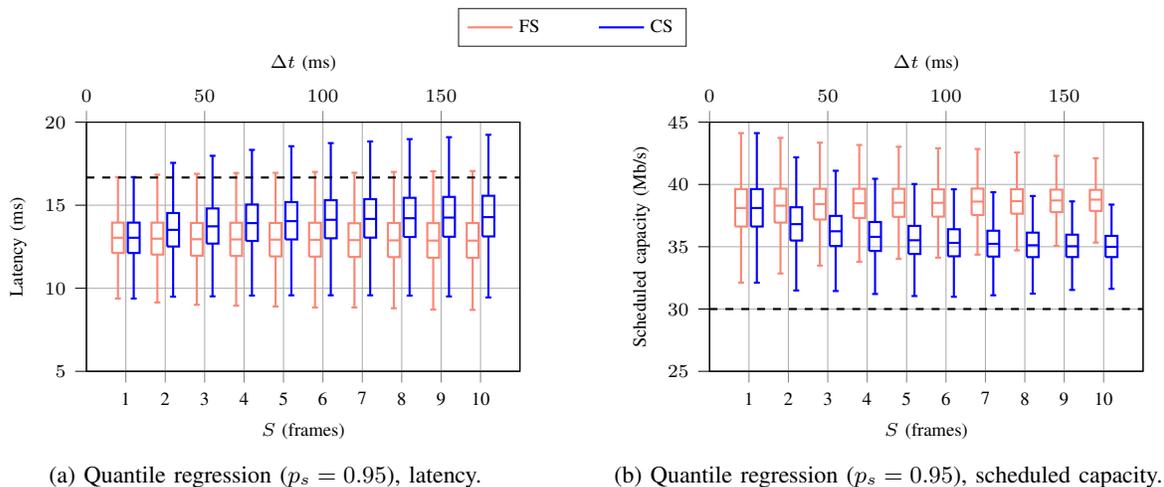
Fig. 9: Boxplot of scheduling performance for FS and CS with $\varphi = 60$ FPS, $R = 30$ Mb/s, $N = 6$, and $\tau = 1$.

(a) Latency.

(b) Scheduled capacity.
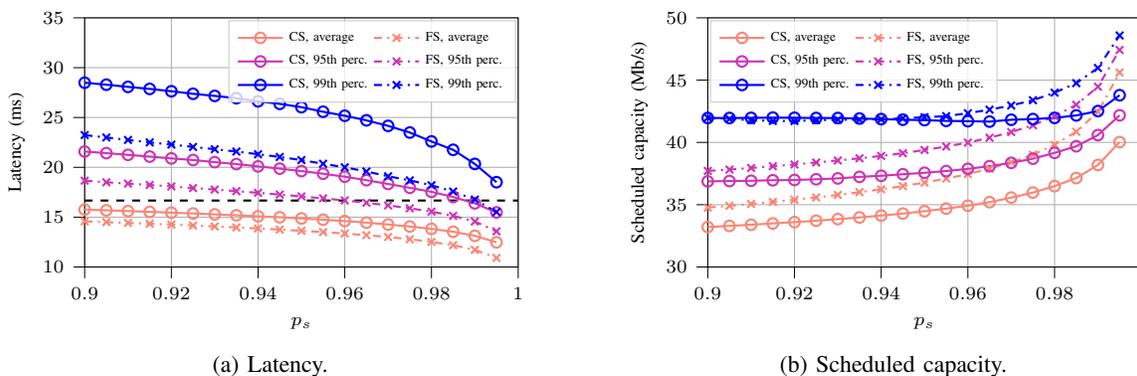
Fig. 10: Average and worst-case percentiles of the latency and scheduled capacity with $S = 6$ (100 ms), $N = 6$, and $\tau = 1$, as a function of the quantile regression parameter p_s.

as S grows, while the average capacity required by the FS algorithm remains roughly constant irrespective of the value of S, but always higher than the capacity used by CS.

This behavior is to be expected, as the errors in frame prediction can compensate over a longer window, but comes at the cost of a higher latency. Naturally, the choice between the two models depends not only on the desired point in the trade-off between QoS and resource efficiency, but also on the capabilities of the underlying system: state-of-the-art slicing frameworks often consider scheduling with a period $\Delta t = 100$ ms, which would correspond to $S = 6$ frames, and the granularity of the scheduling over time and frequency will dictate whether FS is even an option.

It is also possible to simply increase the value of $C(t)$, e.g., by increasing p_s, in the CS scheme to match the FS performance in terms of latency, but CS will always be less efficient for the same latency target. Fig. 10 shows the scheduling performance as a function of the value of p_s. Naturally, a higher p_s means a more conservative prediction of getting larger frames, which reduces the latency but increases the capacity requirements. The closer we get to 1, the more

increasing p_s affects the latency, with a correspondingly larger increase in the capacity that is reserved to the XR flow. We can also notice that CS requires a much higher value of p_s to get the same performance as FS in terms of latency. A sensible example is to target a latency of one inter-frame interval, i.e., $\varphi^{-1} = 16.67$ ms (the dashed line in Fig. 10a), with a probability of 0.95 (the pink lines in Fig. 10). We notice that to meet this requirement, a value of $p_s \geq 0.96$ has to be chosen for the FS scheme, but the same requirement can only be fulfilled if $p_s \geq 0.99$ using CS. This corresponds to an average scheduled capacity of at least 37.45 Mb/s for FS, but 38.18 Mb/s for CS. While the difference is not very significant, and the CS scheme can be used without a big performance loss, choosing the correct value of p_s to compensate for the scheduler's optimism is not simple, particularly in more complex network scenarios, while it is relatively straightforward for FS.

VI. CONCLUSIONS

This work aims at closing a gap in the literature on traffic source modeling: there are several analyses for passive streaming, both 2D and in immersive setups with Head Mounted

Devices (HMDs), and some for live gaming traffic in 2D, but none for interactive XR with strict latency requirements and quasi-CBR encoding. We analyzed live captures from a setup we devised, publishing both the dataset and the code for the analysis, and presented the performance of two regression models. The models are simple and flexible, can be generalized over different traces with limited performance loss, and can be used for provisioning. We also showed a simple Network Slicing (NS) scenario, which highlights the importance of the trade-off between resource efficiency and QoE.

This is a first step towards fully designing an NS system able to satisfy the stringent QoS requirements of XR applications in industrial settings, in which the consequences of network failures are not only discomfort and nausea for the user, but also significant delays in production and even safety hazards. There are several additional analyses and opportunities for future work, that can be divided in two main directions.

The first potential avenue of research is a wider characterization, with different encoding parameters and even different encoders, and considering different applications, going beyond simple VR games to include the industrial and commercial use cases we mentioned above, and a wider set of subjects. The traces should also integrate a record of the head movements of the users, as they correspond to shifts in the point of view of the XR headset and are expected to be strongly correlated with frame size changes.

The other challenge is to actually design schemes and scheduling algorithms able to take into account the nature of the traffic and accommodate it, efficiently exploiting the prediction and adapting to the peculiarities of different communication technologies or even multiple independent links. In this sense, the allocation of resources over time and frequency, the prioritization of users and traffic types, and even the use of packet-level coding to protect the stream from link failures and deep fading events, can be promising avenues to design a solid framework to support XR in mission-critical scenarios. The study of these techniques at all levels of the communication stack, simulating connection impairments in repeatable conditions through a full-stack network simulator, is our first priority in the ongoing work on this subject.

REFERENCES

[1] Huawei Technologies Co., "Empowering consumer-focused immersive VR and AR experiences with mobile broadband," Huawei Technologies Co., White Paper, 2016. [Online]. Available: https://www.huawei.com/en/industry-insights/outlook/mobile-broadband/insights-reports/vr-and-ar

[2] ——, "AR Insight and Application Practice," Huawei Technologies Co., White Paper, 2021. [Online]. Available: https://carrier.huawei.com/~/media/CNBGV2/download/bws2021/ar-insight-and-application-practice-white-paper-en.pdf

[3] Oculus Business, "Virtual Reality – Set to Enter the Business Mainstream," Oculus Business, White paper, Sep. 2020. [Online]. Available: https://go.facebookinc.com/security-whitepaper.html

[4] Qualcomm Technologies, "The Mobile Future of eXtended Reality (XR)," Qualcomm Technologies, Presentation, Nov. 2020. [Online]. Available: https://www.qualcomm.com/media/documents/files/the-mobile-future-of-extended-reality-xr.pdf

[5] 3GPP, "Extended Reality (XR) in 5G," 3GPP, Technical Report (TR) 26.928, Dec. 2020.

[6] ITU-T, "Requirements for mobile edge computing enabled content delivery networks," ITU, Report F.743.10, Nov. 2019.

[7] H. G. Kim, W. J. Baddar, H.-t. Lim, H. Jeong, and Y. M. Ro, "Measurement of exceptional motion in VR video contents for VR sickness assessment using deep convolutional autoencoder," in *ACM Symposium on Virtual Reality Software and Technology (VRST)*, Gothenburg, Sweden, Nov. 2017.

[8] A. A. Barakabitze, A. Ahmad, R. Mijumbi, and A. Hines, "5G network slicing using SDN and NFV: A survey of taxonomy, architectures and future challenges," *Computer Networks*, vol. 167, p. 106984, Feb. 2020.

[9] Nokia, "Cloud gaming and 5G – Realizing the opportunity," Nokia, White Paper, 2020. [Online]. Available: https://onestore.nokia.com/asset/207843

[10] Huawei, "Preparing for a Cloud AR/VR Future," Huawei, White Paper, 2017. [Online]. Available: https://www-file.huawei.com/-/media/corporate/pdf/x-lab/cloud_vr_ar_white_paper_en.pdf?la=en

[11] M. Lecci, A. Zanella, and M. Zorzi, "An ns-3 Implementation of a Bursty Traffic Framework for Virtual Reality Sources," in *ACM Workshop on ns-3 (WNS3)*, Online Event, US, Jun. 2021.

[12] M. Lecci, M. Drago, A. Zanella, and M. Zorzi, "An Open Framework for Analyzing and Modeling XR Network Traffic," *IEEE Access*, vol. 9, pp. 129 782–129 795, 2021.

[13] J. Latta and D. Oberg, "A conceptual virtual reality model," *IEEE CG&A*, vol. 14, no. 1, pp. 23–29, Jan. 1994.

[14] S. Tanwir and H. Perros, "A Survey of VBR Video Traffic Models," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 1778–1802, Fourth Quarter 2013.

[15] S. Liew and D.-Y. Tse, "A control-theoretic approach to adapting VBR compressed video for transport over a CBR communications channel," *IEEE/ACM Trans. on Networking*, vol. 6, no. 1, pp. 42–55, Feb. 1998.

[16] N. Mohsenian, R. Rajagopalan, and C. A. Gonzales, "Single-pass constant- and variable-bit-rate MPEG-2 video compression," *IBM J. Res. Dev*, vol. 43, no. 4, pp. 489–509, Jul. 1999.

[17] M. Carrascosa and B. Bellalta, "Cloud-gaming: Analysis of Google Stadia traffic," in *arXiv*, Sep. 2020.

[18] A. Di Domenico, G. Perna, M. Trevisan, L. Vassio, and D. Giordano, "A Network Analysis on Cloud Gaming: Stadia, GeForce Now and PSNow," *Network*, vol. 1, no. 3, pp. 247–260, Oct. 2021.

[19] P. Graff, X. Marchal, T. Cholez, S. Tuffin, B. Mathieu, and O. Festor, "An Analysis of Cloud Gaming Platforms Behavior under Different Network Constraints," in *17th International Conference on Network and Service Management (CNSM)*, Virtual Conference, Oct. 2021.

[20] O. S. Peñaherrera-Pulla, C. Baena, S. Fortes, E. Baena, and R. Barco, "Measuring Key Quality Indicators in Cloud Gaming: Framework and Assessment Over Wireless Networks," *Sensors*, vol. 21, no. 4, Feb. 2021.

[21] B. Hentschel, M. Wolter, and T. Kuhlen, "Virtual Reality-Based Multi-View Visualization of Time-Dependent Simulation Data," in *IEEE Virtual Reality Conference*, Lafayette, LA, USA, Mar. 2009, pp. 253–254.

[22] F. Chiariotti, "A survey on 360-degree video: Coding, quality of experience and streaming," *Computer Communications*, vol. 177, pp. 133–155, Sep. 2021.

[23] M. Chen, W. Saad, and C. Yin, "Virtual Reality Over Wireless Networks: Quality-of-Service Model and Learning-Based Resource Management," *IEEE Trans. on Communications*, vol. 66, no. 11, pp. 5621–5635, Nov. 2018.

[24] M. Chen, W. Saad, C. Yin, and M. Debbah, "Data Correlation-Aware Resource Management in Wireless Virtual Reality (VR): An Echo State Transfer Learning Approach," *IEEE Trans. on Communications*, vol. 67, no. 6, pp. 4267–4280, Jun. 2019.

[25] X. Yang, Z. Chen, K. Li, Y. Sun, N. Liu, W. Xie, and Y. Zhao, "Communication-Constrained Mobile Edge Computing Systems for Wireless Virtual Reality: Scheduling and Tradeoff," *IEEE Access*, vol. 6, pp. 16 665–16 677, Mar. 2018.

[26] Nvidia, "Video Codec SDK Documentation," 2021. [Online]. Available: https://docs.nvidia.com/video-technologies/video-codec-sdk/nvenc-video-encoder-api-prog-guide/

[27] S. M. Stigler, "Gauss and the invention of least squares," *the Annals of Statistics*, pp. 465–474, May 1981.

[28] R. Koenker and G. Bassett Jr, "Regression quantiles," *Econometrica: journal of the Econometric Society*, pp. 33–50, Jan. 1978.

[29] C. Yu and W. Yao, "Robust linear regression: A review and comparison," *Communications in Statistics-Simulation and Computation*, vol. 46, no. 8, pp. 6261–6282, Sep. 2017.