Demo of `QoEyes`: Towards Virtual Reality Streaming QoE Estimation Entirely in the Data Plane

Francisco Vogt*, Fabricio Rodríguez Cesen*, Ariel Góes de Castro+
Marcelo Caggiani Luizelli+, Christian Esteve Rothenberg*, Gergely Pongrácz|
+Federal University of Pampa (UNIPAMPA), Brazil
*University of Campinas (UNICAMP), Brazil
|Ericsson Research, Hungary

***Abstract*—Recent advances in virtual reality (VR) technology have created new user experiences (e.g., online events, gaming). However, ensuring the user experience is still a challenge. Mostly because Quality of Experience (QoE) measurement is limited to the user or control plane, causing high latencies for different scenarios (e.g., 5G networks and beyond). To address this challenge, we present `QoEyes`, an in-network QoE estimation technique based on Inter-Packet-Gap (IPG) measured in programmable devices. Our results show that a strong estimate of the user's QoE can be provided by measuring the IPG on the data plane. Additionally, in this demonstration, we show this QoE estimate and other related metrics in real time, using a Grafana dashboard running in our monitoring server.**

I. Introduction

Virtual Reality (VR) video streaming applications are no longer a futuristic concept but a present reality. By the end of 2023, the number of VR applications is projected to reach almost 34 million, resulting in at least a 12-fold [1] increase in associated network traffic. This significant growth poses a significant challenge for network operators since VR video streaming applications require high network performance to maintain a reasonable Quality of Experience (QoE). Recent research studies [2] have revealed that VR video applications require a network delay of less than 9ms and can have bandwidth requirements that go beyond 500 Mbps.

To lower the bandwidth requirements of VR video streaming, spherical-to-plane projection is used by VR video streaming players, with the tile-based scheme [3] being a typical example. In this approach, VR videos are encoded at varying resolutions (e.g., 720p, 1080p, or 4K) and then broken into spatial segments, commonly referred to as tiles and temporal segments. While streaming, the VR player will request only those segments and tiles corresponding to the visible area of the complete 360-degree panoramic view, also known as the viewport. Additionally, the VR player utilizes other strategies, such as Adaptive Bitrate and Buffer Management heuristics. This approach allows the player to request higher-quality segments that are predicted to belong to the viewport, while lower resolutions are used for the other tiles.

Recent research efforts have focused on evaluating and estimating the QoE of VR video streaming. However, these studies are usually conducted in the user plane (VR player) or the control plane. Filho et al. [2] proposed a two-stage Machine Learning (ML) assisted approach to infer how users perceive the streamed VR video's performance at the user plane. Despite a few recent initiatives, little has been done to infer QoE in the data plane directly. Analyzing the QoE of each VR video session in the data plane would provide the benefit of reacting to different network conditions in real time.

This demo demonstrates `QoEyes` [4], an in-network QoE estimation approach that utilizes Inter-Packet-Gap (IPG) in programmable network devices. Generally, IPG refers to the difference in arrival time between two consecutive packets of the same network flow. However, it can also be utilized in constructing more robust metrics. By measuring the IPG of each VR video session, `QoEyes` can infer QoE directly on the data plane and make decisions to improve it. We implemented `QoEyes` in P4 using Tofino hardware and evaluated its performance using publicly available VR video traces. Also, we leveraged Grafana [5] visualization tool to provide end-user-oriented demo (see Fig. 1). The results indicate that `QoEyes` can strongly correlate the calculated IPG with the achieved Mean Opinion Score (MOS) and report this in real-time. The demo's main contributions can be summarized as follows:

- demonstration of an in-network VR video streaming QoE estimation directly in the data plane;
- a prototype implementation using Barefoot Tofino HW;
- an open-source software artifact for reproducibility, which can be accessed at https://github.com/intrig-unicamp/QoEyes;
- integration with a Grafana dashboard for real-time visualization;

This work was supported by the Innovation Center, Ericsson S.A. and by the Sao Paulo Research Foundation (FAPESP), grant 2021/00199-8, CPE SMARTNESS. This study was partially funded by CAPES, Brazil - Finance Code 001. This work was partially funded by National Council for Scientific and Technological Development (CNPq 404027/2021-0), Foundation for Research of the State of Sao Paulo (FAPESP 2021/06981-0, 2021/00199-8, 2020/05183-0), and Foundation for Research of the State of Rio Grande do Sul (21/2551-0000688-9).

II. Background and Related Work

A. QoE and VR Video Streaming

QoE measurement can be classified into subjective and objective methods. Subjective methods provide accurate QoE measurements as they measure video quality perceived by the Human Visual System (HVS) [6]. Objective methods are based on quantifiable data, such as bandwidth and latency. To better capture the user's experience, it is best to use both subjective and objective metrics. For VR video streaming, the immersive experience it provides requires high bandwidth, low image quality, and limited interactivity. Therefore, VR video streaming requires decomposing the video into tiles for transmission.

B. Inter-Packet Gap (IPG)

IPG is a network metric that has proven to be efficient in solving several network problems such as microbursts [7] and heavy-hitter [8] detection. The IPG refers to the arrival time difference between two consecutive network packets. This difference can be calculated according to Equation 1, where TS_l and TS_p are the arrival time of the last and penultimate packets, respectively. However, some works explore variations of this metric, for example, calculating the IPG as a Cumulative Sum (CUSUM), Exponentially Weighted Moving Average (EWMA), or Double Exponentially Weighted Moving Average (DEMA) [9]. In this work, we use the IPG calculation as an EWMA in a programmable network device to estimate QoE in real time, enabling a quick response time to a possible loss of QoE.

$$IPG = TS_l - TS_p \tag{1}$$

C. Related Work

The VR-EXP [3] platform enables adaptive tile-based schemes for various network conditions. Meanwhile, PREDICTIVE [2] is a two-stage ML-assisted approach that predicts how users perceive VR video play-out performance based on network Quality of Service (QoS) and tiling schemes. Vidhya et al. [10] use a fuzzy logic mechanism within the Network Data and Application Function (NWDAF) entity in 5G to evaluate QoE and propose an anticipatory scheduling technique to address network delay and congestion. Similarly, Schwarzmann et al. [11] leverages NWDAF standardized interface capabilities in 5G networks to estimate the accuracy of different state-of-the-art regression techniques. Chen et al. [12] offer an Structural Similarity (SSIM)-based approach for assessing the quality of 360-degree videos, which measures structural and textural similarities, while FastInter360 [13] reduces the encoding time of 360-degree videos with EquiRectangular Projection (ERP) by clustering texture blocks into several texture types. Upenik et al. [14] benchmark Peak Signal-to-Noise Ratio (PSNR)-based approaches against ground-truth subjective quality data. Iurian et al. [15] analyze the impact of priority queues in 360-degree video streaming scenarios, evaluating several objective quality metrics and their correlation with subjective quality scores, but do not provide QoE inference. Finally, Bhat et al. [16] infer objective QoE metrics and model QoE value score for VR streaming sessions by leveraging Q-in-Q (IEEE 802.1ad) tunneling and translating the application- into link-layer header information at the edge. Despite previous research efforts to achieve QoE through the data plane, QoEye represents the first approach to design an in-network VR video streaming QoE estimation directly.

III. Inferring QoE in the Data Plane

QoE estimation is typically performed using either measurement or prediction methods on either the user plane or the control plane. While these methods provide accurate results, they may not always be as quick as needed due to the complexity of obtaining certain network metrics and the efficiency of the algorithms. This can limit the prompt implementation of network policies to address potential problems or quality degradation, thus negatively affecting the user QoE.

To address this limitation, the authors propose a novel approach to QoE estimation entirely in the data plane. By assessing QoE in the data plane, decisions can be made at a nanosecond level, thus improving the reaction time to problems and enhancing user QoE. However, QoE inference in the data plane presents a set of challenges, including hardware resource limitations, operational constraints, and limited information about data flow.

In this work, we present QoEyes, a QoE estimation technique that uses IPG calculation to carry out QoE estimation entirely in the data plane. Our approach is based on calculating the IPG as an EWMA (see Equation 2) and exploring the direct correlation between IPG and QoE [9].

$$IPG_w = \alpha \cdot IPG_{w-1} + (1 - \alpha) \cdot IPG_c \tag{2}$$

Our technique is implemented in P4 and calculates the IPGw in the ingress process pipeline. We compute the IPGw per flow and store the result in registers indexed by flow IDs (a hash of src and destination IPs and ports). Note that with this, we can already make decisions and apply network policies based on the calculated IPGw (e.g., use priority queues when detecting a low QoE). However, in this work, we only report the calculated IPGw (QoE estimate) using In-Band Network Telemetry (INT). In the monitoring server, we store the data received in a database using Influxdb and plot it using a dashboard in Grafana. Figure 1 shows an example of this process. More details about our strategy can be obtained in the complete paper [4].

During the demo. In our demonstration, we will present different experiments to attendees, including varied network scenarios. We will run QoEyes remotely in physical Tofino hardware connected with different physical servers (Clients, server, and monitoring server). The main topics that we will present are:

- Experiments as performed in the complete paper, exploring different network scenarios, being able to vary the number of active sessions, network metrics (bandwidth, latency), and background traffic.

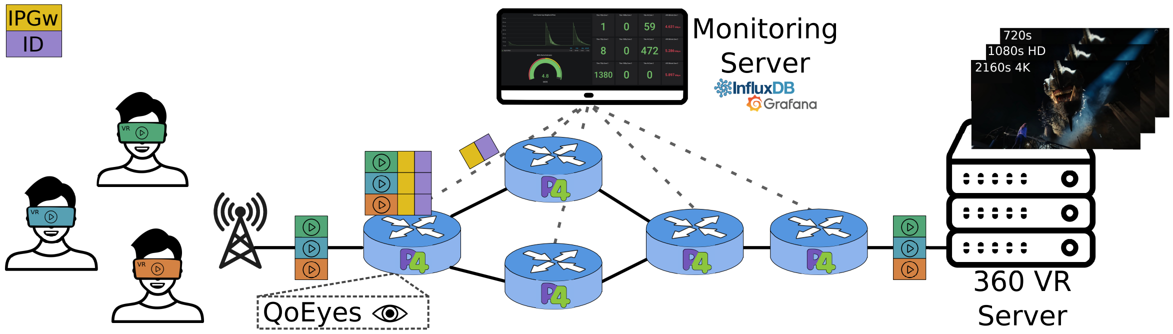

Fig. 1. Illustrating QoEyes architecture in a P4 based network

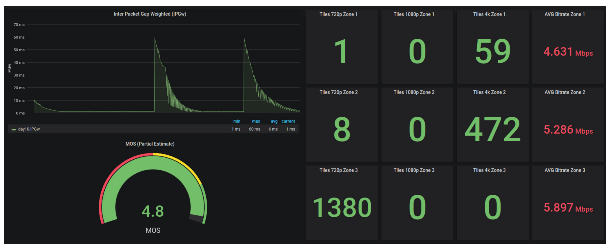

Fig. 2. QoEyes dashboard.

- Attendees will be asked to choose the network conditions (metrics and active traffic) and analyze the real-time impact of these changes on the measured metrics.
- a real-time visualization of our QoE estimation method that will contribute to validating on-the-fly accuracy and show its potential and current limitations.

With these variations in network scenarios (bandwidth, latency, etc.) we will be able to observe how the QoE estimated in the data plane and the QoE measured by the VREXP player behave, demonstrating their correlation. For example, we can see the measured QoE in a congested network scenario, with a low available bandwidth and a high number of active video sections. Figure 2 presents the dashboard that will be used to demonstrate the metrics in real-time. On the dashboard, we can monitor metrics such as the IPGw measured on the switch, and the calculation of the estimated MOS (calculated on the client), demonstrating their correlation. In addition, we can see the advancement of more specific metrics, such as the number of tiles received at 4k, 1080p, and 720p qualities in each zone. For more details about these tiles and their qualities, read the complete paper [4].

IV. Conclusions and Future Work

In this work, we present a demo of QoEyes, a complete data plane technique of QoE estimation for 360-degree (VR) videos. We present a dashboard where we can see the operation of QoEyes in real-time, visualizing the QoE estimate calculated in the data plane and the metrics of the video received by the user. With this, we can compare the effectiveness of our QoE estimation method in real time, modifying the network landscape and visualizing the impact of this on the measured metrics. In future work, we plan to improve our dashboard and data collection, increasing monitoring metrics (e.g., packet size, TCP used sessions) and improving our analysis methods (e.g., thresholds, alerts).

References

[1] Cisco, "Cisco visual networking index: Forecast and trends, 2017-2022," Cisco Systems, Tech. Rep., 2019.

[2] R. I. T. da Costa Filho *et al.*, "Predicting the performance of virtual reality video streaming in mobile networks," in *Proceedings of the 9th ACM Multimedia Systems Conference*, 2018, pp. 270–283.

[3] R. I. T. D. C. Filho *et al.*, "Dissecting the performance of vr video streaming through the vr-exp experimentation platform," *ACM TOMM*, vol. 15, no. 4, pp. 1–23, 2019.

[4] F. G. Vogt, F. Rodriguez Cesen, A. Goes de Castro, M. Caggiani Luizelli, C. Esteve Rothenberg, and G. Pongrácz, "QoEyes: towards virtual reality streaming QoE estimation entirely in the data plane," in *NetSoft*, Madrid, Spain, Jun. 2023.

[5] M. Chakraborty and A. P. Kundan, "Grafana," in *Monitoring Cloud-Native Applications: Lead Agile Operations Confidently Using Open Source Software*. Springer, 2021, pp. 187–240.

[6] J. van der Hooft *et al.*, "Objective and subjective qoe evaluation for adaptive point cloud streaming," in *2020 twelfth international conference on quality of multimedia experience (QoMEX)*. IEEE, 2020, pp. 1–6.

[7] F. G. Vogt, F. Rodriguez, C. Rothenberg, and G. Pongrácz, "Innovative network monitoring techniques through in-band inter packet gap telemetry (ipgnet)," in *Proceedings of the 5th International Workshop on P4 in Europe*, 2022, pp. 53–56.

[8] S. K. Singh *et al.*, "Hh-ipg: Leveraging inter-packet gap metrics in p4 hardware for heavy hitter detection," *IEEE Transactions on Network and Service Management*, 2022.

[9] R. U. Mustafa and C. E. Rothenberg, "Machine learning assisted real-time dash video qoe estimation technique for encrypted traffic," in *Proceedings of the 1st Mile-High Video Conference*, 2022, pp. 123–123.

[10] R. Vidhya, P. Karthik, and S. Jamadagni, "Anticipatory qoe mechanisms for 5g data analytics," in *COMSNETS*. IEEE, 2020, pp. 523–526.

[11] S. Schwarzmann *et al.*, "Ml-based qoe estimation in 5g networks using different regression techniques," *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 3516–3532, 2022.

[12] S. Chen, Y. Zhang, Y. Li, Z. Chen, and Z. Wang, "Spherical structural similarity index for objective omnidirectional video quality assessment," in *IEEE ICME*. IEEE, 2018, pp. 1–6.

[13] I. Storch *et al.*, "Fastinter360: A fast inter mode decision for hevc 360 video coding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 5, pp. 3235–3249, 2021.

[14] E. Upenik, M. Rerabek, and T. Ebrahimi, "On the performance of objective metrics for omnidirectional visual content," in *2017 ninth international conference on quality of multimedia experience (QoMEX)*. IEEE, 2017, pp. 1–6.

[15] C.-M. Iurian *et al.*, "Video streaming evaluation using priority queuing in p4 programmable networks," in *RoEduNet*. IEEE, 2022, pp. 1–5.

[16] D. Bhat *et al.*, "Application-based qoe support with p4 and openflow," in *IEEE INFOCOM*. IEEE, 2019, pp. 817–823.