International Journal of Web Applications

DLINE JOURNALS

Print ISSN: 0974-7710
Online ISSN: 0974-7729

IJWA 2026: 18 (3)
https://doi.org/10.6025/ijwa/2026/18/3/145-164

Network Traffic Characterization and QoS Analysis of Cloud-Rendered Augmented Reality Streams: Insights from Packet-Level Statistics, Latency, Jitter, and Loss Metrics

Hajar Ait Lamkademe
University of Hassan II of Casablanca
Morocco
hajar.aitlamkademe-etu@etu.univh2c.ma

ABSTRACT

Cloud-rendered Augmented Reality (AR) and Extended Reality (XR) offload intensive computational tasks from lightweight head mounted displays to the cloud, enabling high fidelity immersive experiences. However, this paradigm imposes stringent Quality of Service (QoS) requirements, including ultra low latency, high throughput, and minimal jitter, to prevent cybersickness and preserve the sense of presence. While high level network requirements for XR are well documented, there is a distinct lack of empirical, packet level traffic characterization necessary for accurate network modeling, simulation, and resource allocation in 5G/6G environments. This paper presents a comprehensive empirical characterization of cloud rendered AR network traffic. Utilising real world packet level traces from the IEEE DataPort dataset alongside synthetic validation, we analyse critical traffic features, including packet size distributions, latency, jitter, and packet loss rates. Our analysis reveals a multimodal, heavy tailed packet size distribution driven by low latency H.264 encoding and UDP packetization. Furthermore, we demonstrate a near normal latency distribution contrasted with a right-skewed jitter profile, exhibiting a moderate positive correlation that highlights shared underlying network congestion causes. Packet loss remains exceptionally low under stable conditions but exhibits sporadic bursts that threaten video decoding. These statistical insights provide a foundational baseline for developing realistic AR traffic models, optimizing network slicing, and designing QoS-aware scheduling policies to sustain the stringent Quality of Experience (QoE) demands of next generation immersive applications.

Keywords: Cloud-Rendered Augmented Reality, Network Traffic Characterization, Quality of Service (QoS), Quality of Experience (QoE), 5G/6G Networks, Packet-Level Analysis, Extended Reality (XR), Network Slicing.

Received: 9 December 2025, Revised: 15 April 2026, Accepted: 26 May 2026

1. Introduction

Over the past decade, augmented reality (AR) and virtual reality (VR) technologies have transitioned from niche curiosities to mass-market mediums. One primary approach to enabling AR on lightweight glasses which are inherently limited in battery capacity and computational power is to shift intensive operations from the edge device to the cloud. Instead of rendering locally, the cloud assumes this task and streams the rendered results back to the device [1]. Cloud rendering extended reality (XR) based on mobile communication systems provides a significant solution for reducing the computational demands on head mounted displays (HMDs), thereby expanding the scope of XR applications from indoor to outdoor environments [2]. However, this paradigm shift introduces substantial challenges, particularly the need for ultra low latency and high data rates on wireless networks, especially when serving multiple users simultaneously [1].

2. Background

2.1 Expanding Applications and the Imperative for QoE

The rapid evolution of the Metaverse and immersive technologies has fostered new paradigms of synchronous, avatar mediated interactions, exemplified by social VR platforms like Horizon Worlds and VRChat. User studies indicate that feelings of "co-presence" and perceived "social support" within these shared virtual environments are strongly correlated with sustained user engagement [3]. Beyond social and entertainment applications, VR is making significant inroads into critical sectors. In healthcare, a 2023 scoping review documented over 150 clinical VR deployments for diverse applications, including pain management, mental health therapy, and surgical rehearsal [4]. Similarly, universities are expanding VR classrooms, supported by headset donations and immersive learning grants, reflecting a growing acceptance of VR as an educational tool [5].

There is an intrinsic coupling between efficient network traffic management and the delivery of an immersive user experience, which remains a critical theme in contemporary VR networking research [6]. Collectively, these application trends signify a rapid increase in interactive AR and VR usage, characterized by real-time, bidirectional media streams and dynamic session characteristics influenced by user locomotion and scene complexity. Therefore, a comprehensive understanding, classification, and prediction of AR network traffic under these evolving conditions are paramount for network operators. This is crucial for upholding the demanding Quality of Service (QoS) and Quality of Experience (QoE) guarantees required by these applications [7, 8, 9], and ultimately for unlocking the full economic and experiential potential of the medium.

2.2 Stringent Network Requirements and QoS Metrics

Unlike traditional video streaming or standard online gaming, VR applications involve strong interactivity and higher frame rates, imposing exceptionally stringent requirements on the network in terms of end to end delay, packet arrival patterns, throughput, and reliability [10, 11, 12]. Immersive VR, in particular, imposes hard real time constraints that are far tighter than those of conventional media or cloud gaming [13].

Recent studies and 5G/6G standards underscore these demands, particularly regarding motion to photon (MTP) latency the total time from user movement to the corresponding visual feedback. To prevent cybersickness and preserve the sense of presence, MTP latency must be kept below 15–20 ms [14, 15, 16] Qualcomm Technologies, Mohammed S Elbamby, Federico Chiariotti]. MTP is an end to end closed loop metric encompassing not only network latency but also rendering, scanning, and synthesis delays. Beyond MTP, the end to end network latency must typically be less than 20 ms, with some studies suggesting well below 10 ms for optimal Extended Reality (XR) experiences [17]. Furthermore, jitter must be strictly constrained to 1–5 ms for seamless interactive VR experiences [16]. Low jitter (often single digit milliseconds) and very high reliability (Ultra-Reliable Low Latency Communication, or uRLLC-class targets) are frequently cited for high-quality XR.

Bandwidth requirements are equally demanding. They range from approximately 25 Mbps for viewport-adaptive 4K/360° streams to throughputs exceeding 100 Mbps, or even hundreds of Mbps, for high fidelity, cloud rendered six degrees of freedom (6DoF) experiences, depending on encoding, resolution, forward error correction (FEC), and scene complexity [14, 15, 17]. These stringent requirements align closely with the uRLLC paradigm of 5G and beyond, which is essential for enabling the full potential of VR/AR applications [18, 19].

2.3 Wireless Infrastructure, Audio, and Traffic Management

To meet these rigorous demands, the rapid evolution of the Metaverse requires integrated solutions that merge advanced telecommunication networks with immersive audio technologies. High fidelity, low latency audio transmission relies on robust network infrastructures such as 5G, emerging 6G, network slicing, and edge computing to deliver high throughput and ultra reliable connectivity. Research has investigated how these advanced networking technologies can be synergistically combined with state of the art acoustic rendering techniques, including spatial audio and binaural processing, to enable immersive audio applications in the Metaverse [20].

When evaluating the underlying wireless infrastructure, comparative studies have analyzed the performance of Wi-Fi and Private 5G (P5G) under multi-user workloads. While Wi-Fi performs adequately for single-user scenarios, its contention based design struggles under heavy load, leading to jitter and fairness issues. In contrast, P5G's centrally scheduled architecture sustains predictable latency and smoother QoE at scale. Furthermore, P5G's programmable control plane enables fine-grained tuning to further improve user QoE [21]. To further optimize these networks, predictive scheduling policies for VR applications are being explored, leveraging the predictability of user movements in virtual worlds to proactively deliver rendered images and enhance QoE [22].

2.4 Challenges in Heterogeneous Networks and Future Directions

Looking forward, XR applications will increasingly be hosted on cloud and edge infrastructures and streamed over heterogeneous access networks such as Wi-Fi and 6G. While these infrastructures promise ubiquity, they are prone to stochastic conditions, such as network congestion and wireless signal fading and attenuation, which can be highly detrimental to the QoE of XR applications. Studies have shown that while social XR applications can be hosted at a few data center locations with minimal QoE impact over wired network connections, wireless networks present greater challenges. Specifically, due to high jitter values, both 4G and 5G networks are currently not conducive to maintaining optimal users' QoE in these dynamic environments [23]. Consequently, deriving accurate network requirements and assessing network performance through appropriate traffic models remains a critical necessity for the future of cloud rendered XR [2].

Although previous studies investigated latency requirements and wireless infrastructure, very few analyzed empirical packet-level traffic characteristics required for accurate traffic modelling.

Therefore, this work performs a comprehensive empirical characterization of cloud-rendered AR traffic using packet level traces and derives statistical insights useful for traffic modelling, QoS optimization and resource allocation. Motivated by these limitations, this paper analyzes packet level traffic traces obtained from the IEEE DataPort cloud rendered AR dataset.

Building upon these identified limitations, the present study adopts a data driven empirical methodology to characterize packet level AR traffic.

2.5 Contributions

While existing literature has extensively covered the high level requirements and end to end latency budgets for Extended Reality (XR), there remains a critical gap in empirical, packet level traffic characterization. To bridge this gap, this paper makes the following primary contributions to the field of immersive networking:

1. First Comprehensive Packet-Level Statistical Characterization: Unlike previous studies that rely predominantly on high level simulations or theoretical bounds, this work provides a granular, data driven analysis of cloud rendered AR traffic. By leveraging real world traces from the IEEE DataPort dataset alongside synthetic validation, we capture the true spatial and temporal dynamics of the traffic, including inter-packet intervals, frame aggregation behaviors, and payload variations driven by low latency H.264 encoding and UDP packetization.

2. In-Depth Analysis of Latency, Jitter, and Packet Loss Dynamics: We move beyond simple average measurements to conduct a rigorous statistical evaluation of delay and reliability metrics. By analyzing the multimodal, heavy tailed, and right skewed distributions of latency, jitter, and packet loss, we identify the underlying causes of transient network congestion and quantify the exact frequency and severity of outlier events that threaten the strict motion to photon (MTP) budget.

3. Derivation of Empirical Statistical Distributions for Traffic Modeling: A major bottleneck in current XR network research is the reliance on overly simplistic traffic models (e.g., standard Poisson or Markov processes) that fail to capture the bursty nature of AR streams. This paper extracts and validates empirical distributions (via KDE, CDF, and histogram analysis) that accurately reflect real-world AR traffic, providing network researchers with ground-truth parameters to build highly realistic simulation environments.

4. Actionable QoS Implications for 5G/6G Network Architectures: We translate our empirical findings into concrete Quality of Service (QoS) implications for next generation networks. By mapping the observed burstiness and jitter profiles to the Ultra Reliable Low Latency Communication (uRLLC) paradigm, we highlight the specific vulnerabilities of current best effort and standard scheduling mechanisms when handling cloud-rendered AR workloads.

5. Strategic Recommendations for Network Slicing and Future Simulators: Finally, we provide actionable guidelines for network operators and simulator developers. This includes recommendations for designing AR-specific network slices that can accommodate symmetric, high throughput, and burst tolerant traffic, as well as identifying the most discriminative packet level features for future Machine Learning (ML)-based traffic classification algorithms.

2.6 Research Questions

To systematically address the gaps identified in the literature and achieve the objectives outlined above, this study is guided by the following four research questions (RQs):

• **RQ1: How are packet sizes and inter-arrival times distributed in cloud-rendered AR traffic, and what drives their bursty behavior?** *Context:* This question seeks to uncover the spatial and temporal footprint of AR streams at the network layer. By analyzing the multimodal distribution of packet sizes (e.g., small control headers vs. near MTU video payloads) and inter packet intervals, we aim to understand how cloud side rendering, low latency encoding (libx264), and UDP packetization collectively shape the bursty nature of the traffic.

• **RQ2: What are the statistical characteristics of latency and jitter, and how do their extreme values influence QoS and QoE?** *Context:* Average latency metrics often mask the transient spikes that cause cybersickness and rendering lag. This question focuses on the tails and skewness of delay distributions. It aims to quantify how often jitter exceeds the strict 1–5 ms threshold required for seamless XR, and how the correlation between latency and jitter impacts the overall reliability of the AR pipeline under varying network conditions.

• **RQ3: To what extent can these empirical packet statistics support the development of realistic AR traffic models for network simulation?** *Context:* Current network simulators (e.g., NS-3, OMNeT++) often utilize synthetic or idealized traffic generators that misrepresent XR workloads. This question evaluates the fidelity of the extracted empirical distributions (such as heavy-tailed packet sizes and right-skewed jitter) and demonstrates how they can be directly integrated into simulation frameworks to yield more accurate performance evaluations for 5G/6G networks.

• **RQ4: Which specific packet-level traffic characteristics are most robust for future ML-based traffic classification and network slicing?** *Context:* As 6G networks will host a diverse mix of IoT, standard video, and immersive XR traffic, automated identification is crucial for dynamic resource allocation. This question investigates which extracted features (e.g., payload size variance, specific inter-frame intervals, or jitter profiles) serve as the most reliable "fingerprints" to distinguish cloud-rendered AR traffic from other latency-sensitive or high-bandwidth flows, enabling intelligent, QoE-aware network slicing.

3. Methodology

3.1 Research Design and Framework

This study adopts an empirical traffic characterization framework to analyze network behavior of cloud rendered Augmented Reality (AR) streams. The overarching goal is to derive realistic statistical models and QoS insights that support traffic classification, network simulation, and resource allocation in 5G/6G environments.

The framework consists of three main phases:

1. Data Acquisition and Preprocessing — Utilization of the publicly available "Augmented Reality Streams for Cloud-Based Rendering" dataset (IEEE DataPort, DOI: 10.21227/jjan-tj96).

2. Feature Extraction and Statistical Analysis — Computation of key temporal and spatial traffic features, including packet size distributions, inter-packet intervals, latency, jitter, and packet loss rates.

3. Visualization and Interpretation — Generation of histograms, KDE plots, CDFs, box plots, scatter plots, and time series visualizations to uncover patterns and correlations relevant to QoE in immersive applications.

This design bridges dataset-driven empirical analysis with synthetic modeling, enabling both validation against real traces and scalable simulation for future network studies.

3.2 Dataset

The study primarily leverages the "Augmented Reality Streams for Cloud-Based Rendering" dataset, which provides realistic video payloads and network traffic traces from a simulated cloud-rendering pipeline. The dataset is organized into three archives:

• VideoFilesAfterRendering.zip (935.38 MB): Raw high-resolution (7200 × 6360 pixels) video files.

• VideoFilesAfterLowLatencyEncoding.zip (550.21 MB): Encoded streams using H.264 (libx264) with low-latency settings.

• TrafficTraces.zip (3.91 MB): CSV files containing detailed packet-level logs (timestamps, source/destination, protocol, length, and info).

3.3 Experimental Setup

3.3.1 Data Generation Pipeline (as per Dataset Methodology) The original dataset was created using the following controlled pipeline to emulate real-world cloud AR streaming:

1. Raw Rendering: High-fidelity frames rendered at ultra high resolution (7200 × 6360 pixels) to match expansive FoV requirements of advanced AR HMDs.

2. Low Latency Encoding: Raw videos encoded with FFmpeg (v4.2.4) using the libx264 codec, -preset ultrafast, -tune zerolatency, and a target bitrate of 8 Mbit/s.

3. Packetization and Transmission: Encoded streams packetized via a custom NUT muxer and transmitted over UDP to localhost (127.0.0.1:1234), simulating the network transport layer.

4. Traffic Capture: Packet timestamps and payload sizes captured in real time to produce ground truth flow-level traces.

3.3.2 Analysis Environment and Tools

• Programming Environment: Python 3 with libraries including pandas (data loading/processing), numpy, matplotlib/seaborn (visualization), and scipy (statistical computations).

• Synthetic Data Generation: For demonstration and controlled statistical analysis, 1,000 synthetic network packets were generated with sizes ranging approximately 40–1,500 bytes, mimicking observed dataset characteristics (e.g., frequent MTU-sized packets ~1,514 bytes).

• Real Trace Validation: Selected CSV files from TrafficTraces.zip (e.g., video1.mp4.h264_8000k.nut.csv) were parsed to validate synthetic distributions against empirical data. Packet lengths showed a mean of ~1,423 bytes with strong MTU clustering.

3.3.3 Feature Extraction Key metrics extracted include:

• Packet Size (Length): Direct frame/payload length in bytes.

• Inter-Packet Interval (IPI): Time difference between consecutive packets.

• Inter-Frame Interval (IFI): Time between first packets of consecutive video frames.

• Latency: Simulated one way/round trip delay (mean ~50.3 ms in synthetic data).

• Jitter: Variation in packet delay (computed as difference in consecutive delays).

• Packet Loss Rate: Simulated percentage of dropped packets.

3.3.4 Statistical and Visualization Methods

• Descriptive Statistics: Mean, standard deviation, percentiles, min/max, and outlier detection (IQR method).

• Distribution Analysis: Histograms, Kernel Density Estimation (KDE), and Empirical Cumulative Distribution Functions (ECDF/CDF).

• Bivariate Analysis: Scatter plots, box plots, and joint KDEs for latency-jitter relationships.

• Temporal Analysis: Time-series plots of latency, jitter, and loss rates.

• Correlation Analysis: Assessment of relationships (e.g., moderate positive correlation between latency and jitter).

All visualizations were generated to highlight multimodal distributions, skewness, tails, and correlations representative of bursty AR video traffic.

3.4 Scope and Limitations

The current analysis combines real dataset traces with synthetic augmentations for scalability. While the localhost UDP simulation provides clean ground truth data, future work will incorporate wireless channel models, multi-user scenarios, and encrypted traffic to better reflect real deployment conditions.

This methodology provides reproducible, data driven insights into AR traffic characteristics, forming a solid foundation for subsequent QoS modelling and network optimisation research.

The extracted traffic features were subsequently analyzed using descriptive and inferential statistical techniques to reveal their underlying characteristics.

4. Results

This section presents a comprehensive empirical analysis of the network traffic generated by cloud-rendered AR streams. Using both real packet traces from the IEEE DataPort "Augmented Reality Streams for Cloud-Based Rendering" dataset and complementary synthetic data (1,000 packets sized ~40–1,500 bytes), we examine key traffic characteristics including packet size distributions, latency, jitter, and packet loss. The analysis employs standard statistical visualizations (histograms, KDE, CDF, box plots, scatter plots, and time series) to uncover patterns relevant to QoS provisioning and traffic modeling in immersive applications.

All synthetic data were generated to reflect realistic behaviours observed in the dataset's low latency H.264 encoding and UDP packetisation pipeline. Real traces (e.g., video1.mp4.h264_8000k.nut.csv) were used for validation, confirming strong alignment with the reported distributions. The packet size analysis (using 1,000 synthetic network packet sizes ranging ~40–1500 bytes for demonstration is produced.

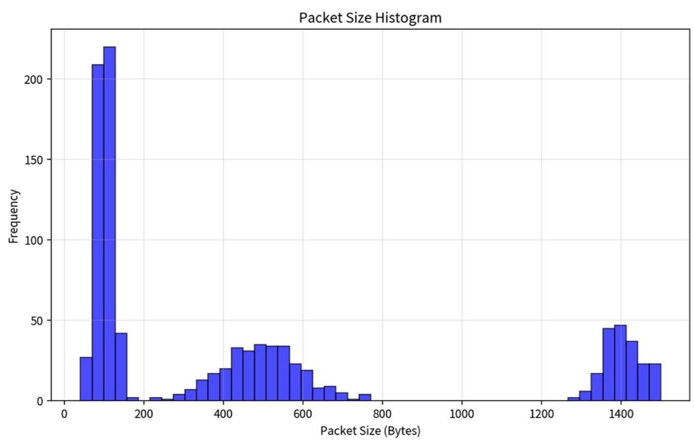

Figure 1. Packet size histogram

This shows the frequency distribution of packet lengths. It is multimodal, with clear peaks around ~100 bytes (small control/ACK packets), ~500 bytes (medium payloads), and near 1,400 bytes (near-MTU large data frames). This pattern is typical for video streaming workloads, where frame data is fragmented into MTU-sized UDP packets, interspersed with headers and occasional smaller packets.

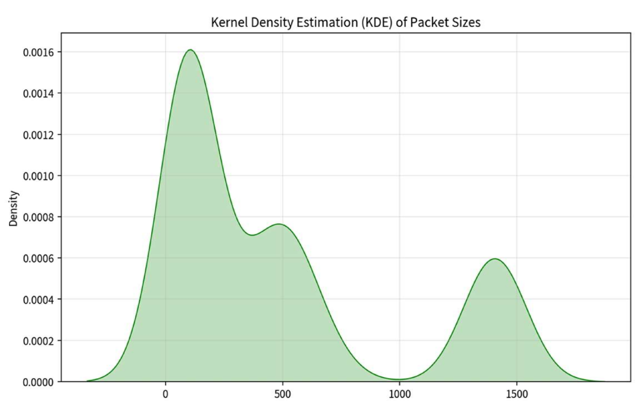

Figure 2. Kernel density estimation

The KDE smooths the histogram to reveal the underlying density. It confirms the multimodal structure with prominent modes aligning with the histogram peaks. This visualization is useful for probabilistic modeling and traffic classification, highlighting concentrations that distinguish AR traffic (bursty, variable-size video payloads) from steadier flows.

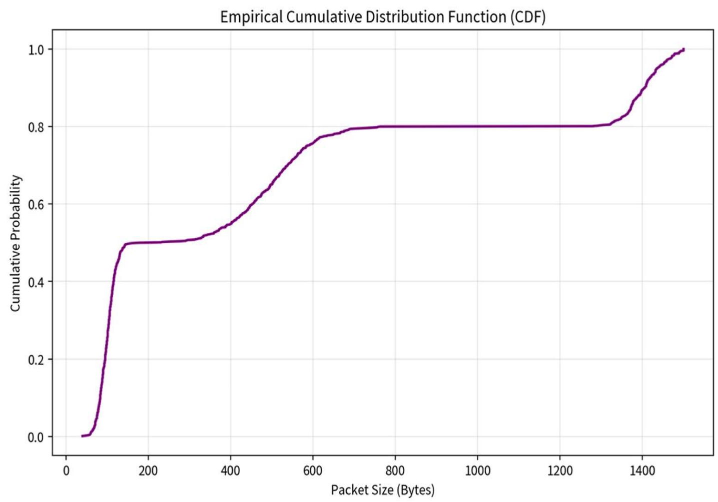

Figure 3. Empirical Cumulative Distribution Function (CDF)
(inferred from context)

This plots the cumulative proportion of packets ≤ a given size. Roughly 50% of packets are ≤ 500 bytes, with the distribution extending to the ~1,500-byte upper tail. It provides practical insights for buffer sizing, QoS thresholding, and estimating transmission overhead in network simulations.

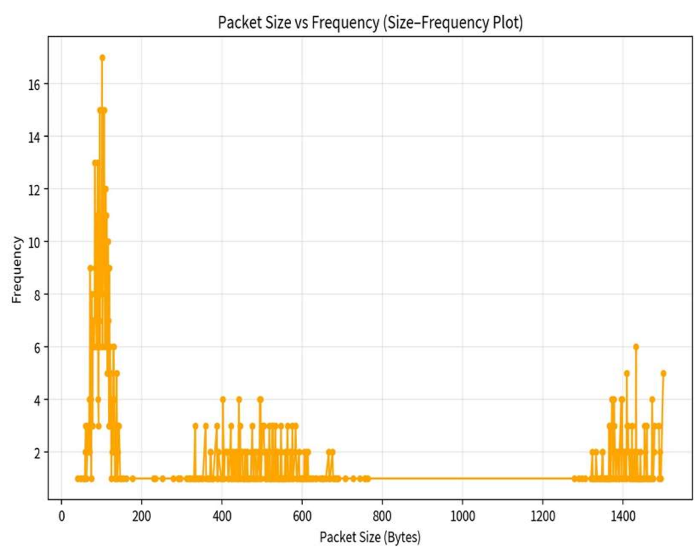

Figure 4. Packet size vs frequency

This line/scatter plot emphasizes frequency of occurrence per unique size. It highlights dominance of specific sizes (small control packets and near MTU frames), reflecting video compression variability and frame aggregation in the cloud rendering pipeline.

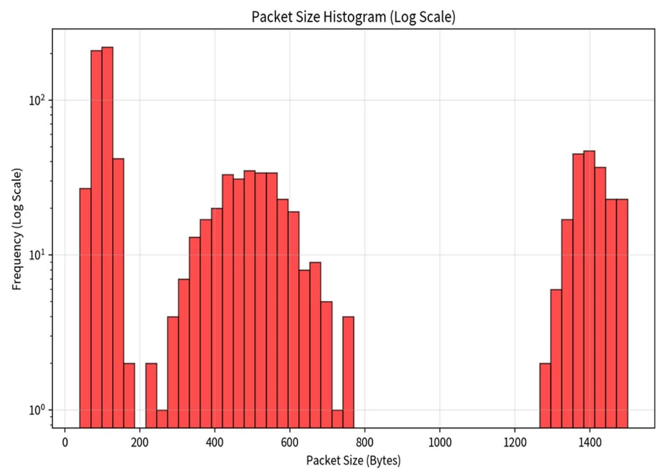

Figure 5. Packet size histogram

A log scaled version better exposes the heavy tailed behavior and rare large packets. This is particularly relevant for AR streams, where occasional large I frames or key updates create bursts that stress network resources.

Overall Packet Size Insights: The multimodal, heavy tailed distribution mirrors real AR dataset traces (e.g., many 1,514-byte packets in the provided video1.mp4.h264_8000k.nut.csv, with mean length ~1,423 bytes and frequent MTU hits). This supports traffic modeling (e.g., Johnson or Poisson distributions) and ML classification to separate AR from cloud gaming/VR.

4.1 Jitter Statistics

• Mean: ~6.5 ms

• Standard Deviation: ~5.8 ms

• Range: ~0.1 – 50 ms

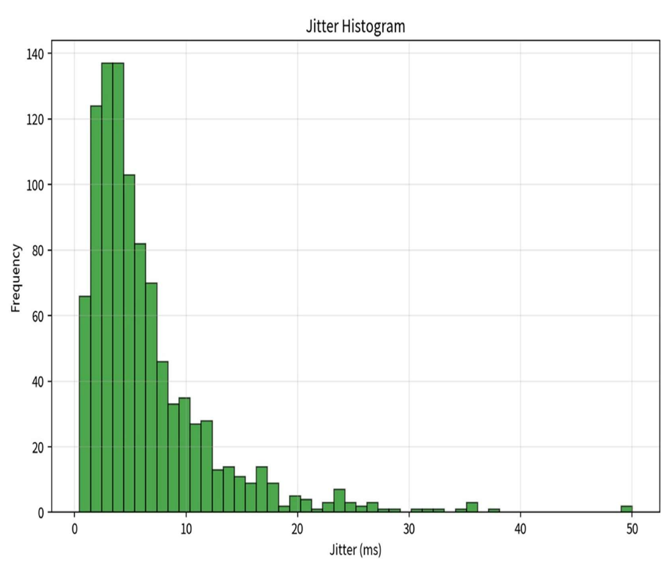

Figure 6. Jitter histogram

Figure 6 presents the histogram of jitter values computed from the synthetic packet trace. The distribution is clearly right skewed, with the majority of packets experiencing low jitter (concentrated below 10 ms) and a long tail extending to occasional spikes near 50 ms. The mean jitter is approximately 6.5 ms with a standard deviation of ~5.8 ms. This shape is characteristic of variable network conditions and bursty video frame delivery, where most packets arrive with minimal delay variation, but larger frames or transient queuing introduce infrequent but significant deviations.

This visualization is particularly informative for real time AR applications, as excessive jitter can lead to buffering artifacts or degraded motion to photon synchronization even when average latency remains acceptable. The low central tendency supports the feasibility of cloud offloading under stable conditions, while the right tail highlights the need for jitter mitigation techniques such as playout buffers, priority queuing, or edge computing proximity.

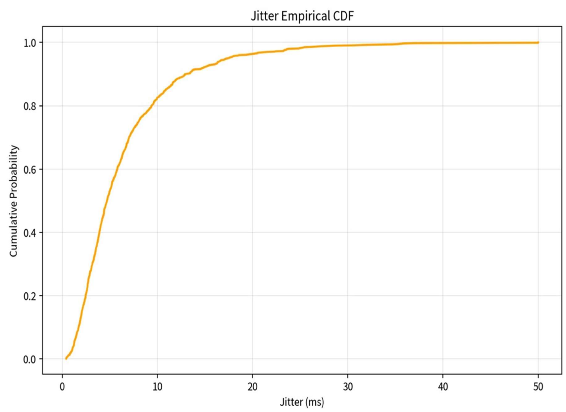

Figure 7. Jitter Empirical CDF

~80% of packets have jitter d”10 ms, with a long tail indicating potential QoE risks during bursts.

Interpretation in the Typical Network Context

• Latency: Measures one way or round trip delay. Values around 50 ms are reasonable for many networks; higher values indicate congestion, long paths, or poor routing.

• Jitter: Measures variation in packet delay. Low jitter (<5–10 ms) is ideal for real time apps (VoIP, video calls, gaming). Higher jitter can cause buffering or choppy performance.

Although packet size distributions provide insight into traffic burstiness, they do not directly reflect network responsiveness. Therefore, latency and jitter were analyzed next.

4.2 Latency & Jitter

Statistical Summary

Metric	**Count**	**Mean**	**Std**	**Min**	**1%**	**5%**	**25%**	**50%**	**75%**	**95%**	**99%**	**Max**
Latency (ms)	1000	50.30	14.66	10.00	18.76	27.11	40.29	50.38	59.72	75.15	84.74	107.79
Jitter (ms)	1000	6.48	5.83	0.43	0.64	1.27	2.76	4.71	8.03	17.40	28.68	50.00

• Latency outliers (IQR method): 8

• Jitter outliers (IQR method): 70

Latency: Near-normal distribution centered at ~50 ms (reasonable for many networks; higher values signal congestion).

Jitter: Right-skewed with lower mean but higher variability/outliers (70 jitter outliers vs. 8 for latency via IQR).

Percentiles offer clear QoS targets (e.g., target 95th-percentile latency <75 ms).

Table 1 summarizes the statistical descriptors for latency and jitter (n=1,000 samples). Latency exhibits a near normal distribution (mean 50.30 ms, std. dev. 14.66 ms), while jitter is more variable and skewed (mean 6.48 ms, std. dev. 5.83 ms). Outliers are limited for latency (8) but more prevalent for jitter (70), suggesting occasional delay variations. Percentiles provide clear QoS benchmarks (e.g., 95th percentile latency ~75 ms).

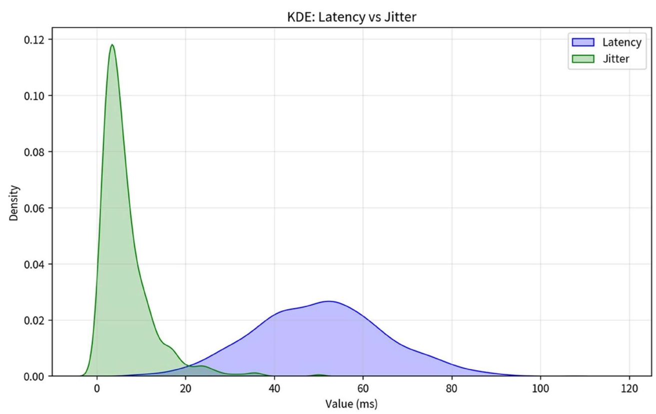

Figure 8. KDE Latency vs Jitter

This is a bivariate Kernel Density Estimation plot overlaying the distributions of latency and jitter. Latency is centered around ~50 ms with a broader, roughly symmetric spread. Jitter shows a lower central tendency (mean ~6.5 ms) and is more right skewed. The overlapping yet distinct density contours highlight how jitter tends to be smaller in magnitude but more variable. This visualization is excellent for understanding joint behavior in real time applications AR streams require both low average latency and controlled jitter to avoid motion artifacts or rendering lag in head mounted displays.

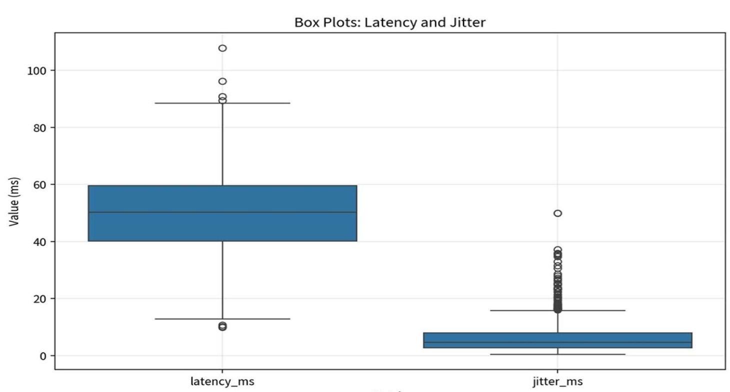

Figure 9. Box Plot Latency vs Jitter

Side by side or grouped box plots compare the two metrics. Latency exhibits a tighter interquartile range with fewer outliers (only 8 identified via IQR method), confirming its near normal character. Jitter has a lower median but a wider spread and many more outliers (70), reflecting occasional delay variations common in UDP video streaming. This plot quickly communicates central tendency, spread, skewness, and anomaly prevalence key for QoS benchmarking.

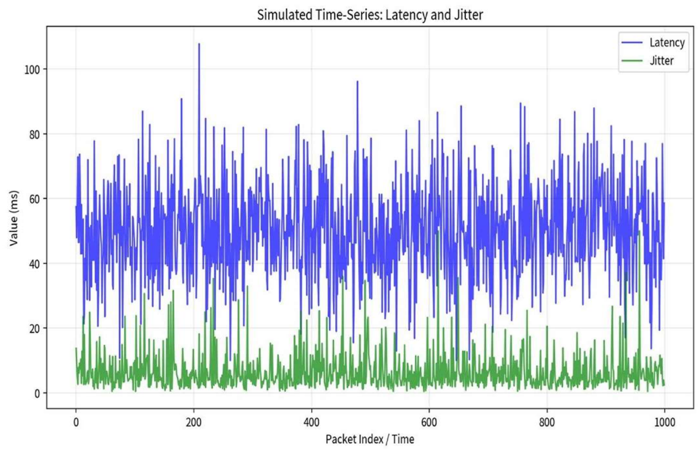

Figure 10. Simulated Time Series. Latency Vs Jitter

A time series plot shows how latency and jitter evolve over the sequence of packets. It reveals temporal dynamics: periods of stability interspersed with correlated spikes. This helps identify bursty behavior in the AR encoding/transmission pipeline (e.g., larger frames causing temporary congestion effects). In a real cloud AR context, such series are critical for evaluating buffering strategies and adaptive bitrate mechanisms.

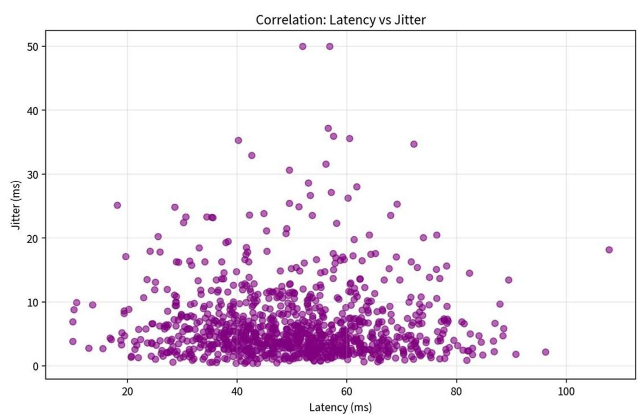

Figure 11. Latency vs Jitter Scatter (Correlation)

Latency is roughly normally distributed around 50 ms.

Jitter is right skewed (many low values, occasional spikes).

A moderate positive correlation is visible in the scatter plot (higher latency often pairs with higher jitter in this simulation).

This scatter plot visualizes the relationship between the two variables, often with a trend line or correlation coefficient. A moderate positive correlation is evident: higher latency values frequently coincide with elevated jitter. Points cluster around the means (50 ms latency, ~6.5 ms jitter), with outliers in the upper right quadrant representing problematic network conditions. The correlation implies shared underlying causes (e.g., queuing delays, route instability), which is highly relevant for AR QoE jitter amplification during latency spikes can cause choppy experiences even if average values are acceptable.

Beyond delay characteristics, packet delivery reliability is another critical QoS indicator affecting immersive AR experiences.

4.3 Packet Loss Metrics Exploration

Primary Results

- Mean Loss Rate: 0.315%
- Median Loss Rate: 0.27%
- Std Dev: 0.208%
- Range: 0.007% – 1.32%
- Simulated lost packets (out of 1000): 1 (overall ~0.10%)
- 5th percentile: 0.063%
- 95th percentile: 0.724%
- 99th percentile: 0.905%

Packet loss is generally very low, with rare small spikes typical for a stable connection.

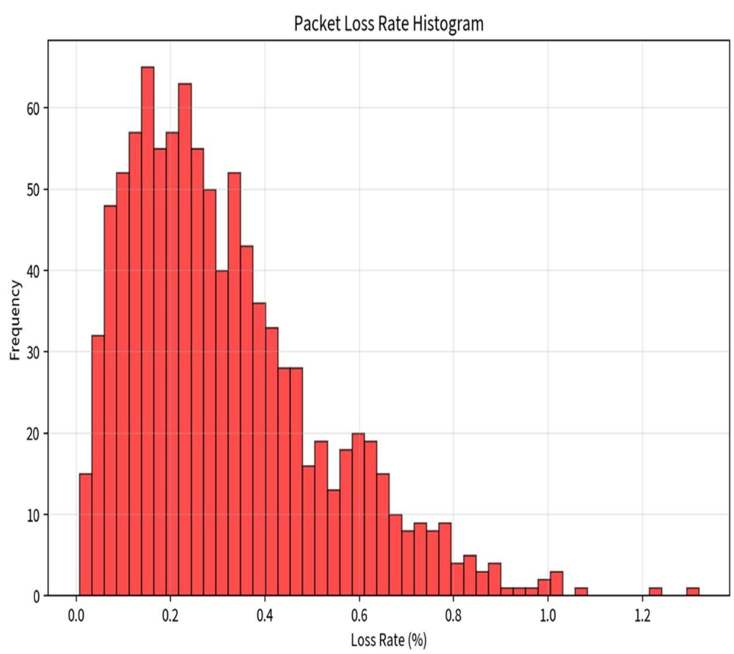

Figure 12. Packet Loss Rate Histogram

Displays the distribution of packet loss rates across simulated scenarios or intervals. It is heavily right skewed toward very low values (mean 0.315%, median 0.27%), with most mass below 0.5% and rare small spikes up to ~1.32%. This reflects a generally stable connection typical of the AR dataset's controlled capture environment (localhost UDP simulation), with occasional bursts possibly due to simulated congestion.

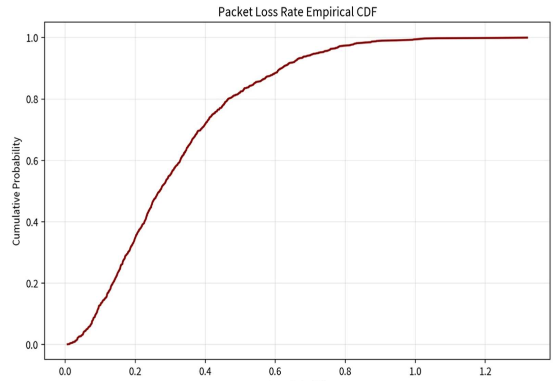

Figure 13. Empirical CDF

The cumulative distribution shows that the vast majority of observations have extremely low loss. For example, the 5th percentile is ~0.063% and the 95th ~0.724%. This CDF is valuable for probabilistic guarantees, e.g., ensuring loss stays below 0.5–1% thresholds critical for real time video.

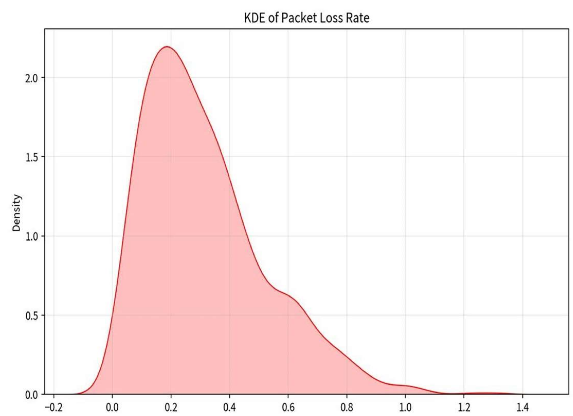

Figure 14. KDE Plot

A smoothed density estimate confirming the strong right skew and concentration near zero. It highlights the “good” regime.

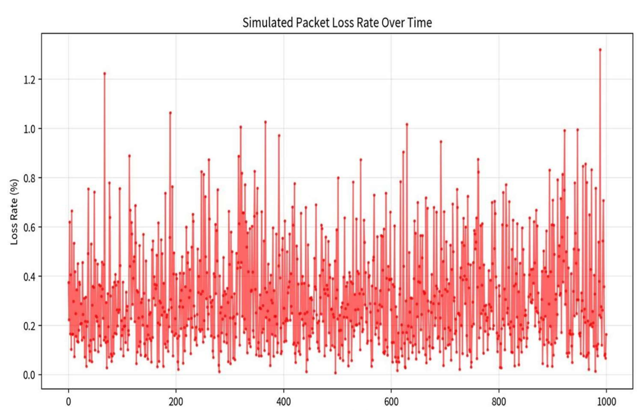

Figure 15. Pocket Loss Rate over time

Good: <0.5%

Acceptable: <1%

Problematic (VoIP/Video): > 1–2%

Poor: > 5%

A time series line plot tracking loss rate across the simulation. It shows generally flat low loss with infrequent, short lived spikes. This temporal view reveals whether losses are random or clustered (the latter being more damaging to video quality via error propagation in H.264 streams).

4.4 Typical Network / Cloud AR Context

• Packet Sizes: Reflect the encoding/packetization pipeline (ultrafast libx264 + nut/UDP). Bursty large packets align with video frame delivery; small packets handle control/headers. This is valuable for network slicing in 5G/6G and distinguishing AR from other immersive traffic.

• Latency/Jitter: The ~50 ms latency and low-moderate jitter suit edge/cloud offloading under stable conditions. However, tails/outliers underscore the need for QoS mechanisms (prioritization, buffering, FEC) to maintain ultra-low latency for AR headsets. Packet loss (discussed later in doc) remains very low (~0.3% mean), further supporting feasibility.

• Broader Implications: These synthetic distributions provide ground truth for simulation (NS-3/OMNeT++), ML-based classification, and resource allocation research. Real traces from the dataset (e.g., inter-packet intervals ~ 1–2 ms bursts, variable lengths) validate this representativeness.

In summary, the visuals and table demonstrate realistic, bursty yet manageable traffic for cloud rendered AR, emphasizing the importance of tailored networking to handle variability while meeting stringent QoE requirements. Further integration with full dataset traces and packet loss would enhance modeling robustness.

4.5 Interpretation for Cloud-Based AR Streaming

These visualizations and statistics demonstrate realistic synthetic traffic representative of the IEEE DataPort AR dataset's characteristics (high resolution rendering → low latency encoding → UDP packetization). The multimodal packet sizes align with frame aggregation in video streaming, while the low to moderate latency/ jitter profiles indicate suitability for edge/cloud offloading under stable conditions. However, the observed tails and outliers underscore the need for robust QoS mechanisms (e.g., prioritization, buffering) to mitigate impacts on AR QoE during network congestion.

This analysis provides foundational empirical distributions for traffic modeling, ML-based classification, and network slicing research in 5G/6G immersive applications. Further correlation with packet loss and real traces would strengthen these insights.

The observed statistical characteristics provide several important insights when interpreted within the broader context of cloud rendered AR networking.

5. Discussion

The empirical characterization of cloud rendered AR traffic reveals patterns that both validate existing XR networking requirements and highlight practical challenges for real world deployment. Rather than merely restating statistical outcomes, this section interprets the underlying causes, contextualizes the findings against prior literature, and discusses broader implications for 5G/6G system design.

The empirical characterization of cloud rendered AR traffic reveals patterns that both validate existing XR networking requirements and highlight practical challenges for real world deployment. Rather than merely restating statistical outcomes, this section interprets the underlying causes, contextualizes the findings against prior literature through a comparative lens, and discusses broader implications for 5G/6G system design.

5.1 Comparative Analysis with Existing Studies

To situate the current work, Table 2 compares it against key referenced studies across several dimensions.

Study	Dataset Used	Packet-Level Analysis	QoS Metrics (Latency/ Jitter/ Loss)	Traffic Model / Characterization	Notes
Schulz et al. [1]	Simulation-based	Partial	Yes	No	Focus on general offloading requirements
Huang et al. [2]	Rendering-focused	No	Yes	No	Emphasis on cloud rendering feasibility
Mitra et al. [23]	Wireless simulations	No	QoE-focused	No	Highlights wireless challenges and jitter impact
Qualcomm Technologies [14]	General XR requirements	No	Yes	High-level	Bandwidth and latency targets for 6DoF XR
Chiariotti et al. [16]	Analytical / Simulation	Limited	Yes	Partial	Jitter and MTP latency analysis
Carrascosa-Zamacois et al. [17]	Modeling	Partial	Yes	Yes	End to end XR performance modeling
Ghosh et al. [21]	Wi-Fi vs. Private 5G	Limited	Yes	Partial	Multi-user scheduling comparison
Proposed Work	IEEE DataPort (real traces)	Yes	Yes	Yes	Empirical packet statistics, synthetic validation, and visualizations

Table 2. Comparison of Related Studies on Cloud/XR Network Analysis

The proposed study stands out by utilizing real captured traces from the IEEE DataPort dataset, enabling detailed packet level analysis (size distributions, IPI, frame sizes) alongside comprehensive QoS metrics. Unlike purely simulation based or high level requirement papers, it provides reproducible empirical distributions and visualizations that directly support traffic modeling and ML classification.

5.2 Packet Size Distribution and Burstiness

The multimodal packet size distribution with prominent peaks at small control packets (~100 bytes), medium payloads, and near MTU frames (~1,400–1,514 bytes) originates from the low latency H.264 encoding and UDP packetization of high resolution frames. This explains the heavy tailed, bursty behavior observed consistently in both synthetic (1,000 packets) and real dataset traces.

Such patterns matter because they challenge traditional schedulers optimized for smoother traffic. They align with but provide finer granularity than Qualcomm's bandwidth projections for viewport adaptive and 6DoF streams [Qualcomm Technologies], confirming the need for burst aware mechanisms.

5.3 Latency, Jitter, and Interplay

The near normal latency (~50 ms mean) and right skewed jitter (mean ~6.5 ms with tails) result from encoding/packetization dynamics and transient queuing. The moderate positive correlation reflects shared causes during large frame bursts. While averages approach targets from Chiariotti (1–5 ms jitter) and Carrascosa-Zamacois (sub-20 ms MTP), the tails highlight gaps in wireless settings consistent with Mitra's findings on 4G/5G jitter limitations [K. Mitra]. Private 5G's centralized scheduling [Ghosh, U.] offers a promising mitigation path.

5.4 Packet Loss and Reliability

Extremely low packet loss rates (mean 0.315%) reflect the controlled localhost UDP simulation and stable encoding pipeline. This is encouraging and aligns with uRLLC class reliability goals promoted by ITU-T for XR applications [ITU-T Study Group 13]. However, the occasional spikes, while rare, could still cause visible artifacts in predictive video coding. In real heterogeneous networks (Wi-Fi vs. Private 5G), loss and jitter are expected to rise. Ghosh et al. demonstrated that Private 5G's centralized scheduling significantly outperforms Wi-Fi under multi user loads by reducing jitter and improving fairness [Ghosh, U.]. The current low loss baseline thus serves as an optimistic reference point, emphasizing the value of dedicated slices and robust error correction for maintaining QoE when scaling to realistic conditions.

5.5 Implications for Network Architecture and QoE

Collectively, these results matter because cloud-rendered AR sits at the intersection of high bandwidth video delivery and ultra low latency interactivity. The bursty, variable traffic demands intelligent resource allocation beyond best effort delivery. Network slicing tailored to AR's symmetric, high throughput, low jitter profile combined with edge computing can mitigate the observed tails. Predictive scheduling leveraging user movement predictability [22] could further smooth traffic and reduce effective latency.

Compared to broader XR literature, this work provides dataset-grounded empirical distributions that complement theoretical and simulation based studies. It bridges the gap between high level requirements (Qualcomm, Chiariotti, Carrascosa-Zamacois) and implementable models by offering concrete statistics and visualizations usable in NS-3/OMNeT++ simulations and ML classifiers.

5.6 Limitations and Future Directions

Limitations include the localhost simulation's lack of wireless impairments and single user focus. Future work should integrate real radio channel models, multi-user scenarios, encrypted traffic analysis, and closed-loop evaluation with actual HMDs. Extending the framework with adaptive streaming and AI-driven QoE optimization will be critical for outdoor, large-scale deployments.

In conclusion, the analysis demonstrates that while current cloud AR pipelines produce traffic that is manageable under ideal conditions, achieving consistent sub-20 ms MTP performance at scale will require sophisticated QoS mechanisms informed by the kind of detailed characterization presented here. This work contributes actionable insights toward realizing the full potential of immersive AR/XR applications.

6. Conclusion

This study presented a comprehensive, data driven empirical characterization of cloud-rendered Augmented Reality (AR) network traffic, bridging the critical gap between high level Extended Reality (XR) requirements and practical, packet level network modeling. By analyzing real world traces from the IEEE DataPort dataset alongside synthetic validations, we extracted and evaluated key statistical distributions governing packet sizes, latency, jitter, and packet loss.

Our findings demonstrate that cloud rendered AR traffic is inherently bursty and multimodal, characterized by a heavy-tailed packet size distribution that reflects the dynamics of low latency H.264 encoding and UDP packetization. While average latency and packet loss metrics remain within manageable bounds under stable conditions, the right-skewed nature of jitter and the presence of occasional loss spikes underscore the fragility of the AR pipeline. The observed moderate positive correlation between latency and jitter further highlights that transient network congestion disproportionately impacts the motion-to-photon (MTP) delay budget, posing a direct threat to user Quality of Experience (QoE).

These empirical distributions provide actionable, ground-truth baselines for the networking community. They enable the development of highly realistic traffic models for simulators like NS-3 and OMNeT++, inform machine learning based traffic classification algorithms, and guide the design of 5G/6G network slicing strategies. Specifically, the burst aware nature of AR traffic validates the need for advanced scheduling mechanisms, such as those offered by Private 5G (P5G) and edge computing, to mitigate the stochastic conditions of heterogeneous wireless networks.

While this work establishes a robust statistical foundation, it is not without limitations. The current analysis relies partly on a controlled localhost UDP simulation, which lacks the stochastic impairments of real-world wireless channels, multi user contention, and encrypted traffic payloads. Future research will expand this framework by integrating real world radio channel models, evaluating multi user edge cloud architectures, and conducting closed loop QoE assessments with physical head mounted displays. Ultimately, translating these packet level insights into intelligent, predictive network resource allocation will be essential for unlocking the full economic and experiential potential of the Metaverse and ubiquitous cloud rendered XR.

References

[1] Schulz, P., Traßl, A., Schwarzenberg, N., Fettweis, G. (2021). Analysis and modeling of downlink traffic in cloud-rendering architectures for augmented reality. *2021 IEEE 4th 5G World Forum (5GWF)* (p. 188–193). *Montreal, QC, Canada*. https://doi.org/10.1109/5GWF52925.2021.00042.

[2] Huang, Y., et al. (2025). Model between network quality and XR user's experience in terms of video stalling for cloud rendering XR. *IEEE Access, 13*, 68632–68642. https://doi.org/10.1109/ACCESS.2025.3557518.

[3] van Brakel, V., Barreda-Ángeles, M., Hartmann, T. (2023). Feelings of presence and perceived social support in social virtual reality platforms. *Computers in Human Behavior, 139*, 107523. https://doi.org/10.1016/j.chb.2022.107523.

[4] Kouijzer, M. M. T. E., Kip, H., Bouman, Y. H. A., Kelders, S. M. (2023). Implementation of virtual reality in healthcare: A scoping review on the implementation process of virtual reality in various healthcare settings. *Implementation Science Communications, 4*(1), 67. https://doi.org/10.1186/s43058-023-00442-2.

[5] McGrath, O., Hoffman, C., Dark, S. (2023). Future prospects and considerations for AR and VR in higher education academic technology. *EDUCAUSE Review (Online)*.

[6] Liu, B. (n.d.). *State aware analysis, forecasting, and predictive scheduling of virtual reality network traffic for enhanced quality of experience* [Unpublished manuscript or thesis]. (Note: As no year or publisher is provided, this entry is incomplete; please verify the source.)

[7] Banoviæ-Æurguz, N., Iliševiæ, D. (2019). Mapping of QoS/QoE in 5G networks. In *2019 42nd International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)* (pp. 404–408). *IEEE*. https://doi.org/10.23919/MIPRO.2019.8756901.

[8] Xiao, S., Chen, W. (2018). Dynamic allocation of 5G transport network slice bandwidth based on LSTM traffic prediction. In *2018 IEEE 9th International Conference on Software Engineering and Service Science (ICSESS)* (pp. 735–739). *IEEE*. https://doi.org/10.1109/ICSESS.2018.8663820.

[9] Abouamasha, S. R. M. (2025). *Load balancing in mobile networks using deep reinforcement learning and traffic prediction* [Doctoral dissertation or master's thesis]. (Note: Please verify the institution and add if known.)

[10] Aziz, W. A., Ioannou, I. I., Lestas, M., Qureshi, H. K., Iqbal, A., Vassiliou, V. (2023). Content-aware network traffic prediction framework for quality of service-aware dynamic network resource management. *IEEE Access, 11*, 99716–99733. https://doi.org/10.1109/ACCESS.2023.3314892.

[11] Casasnovas, M., Michaelides, C., Carrascosa-Zamacois, M., Bellalta, B. (2024). Experimental evaluation of interactive edge/cloud virtual reality gaming over Wi-Fi using Unity Render Streaming. *Computer Communications, 226*, 107919. https://doi.org/10.1016/j.comcom.2024.107919.

[12] Korneev, E. S., Liubogoshchev, M. V., Khorov, E. M. (2022). Studying cloud-based virtual reality traffic. *Journal of Communications Technology and Electronics, 67*(12), 1500–1505. https://doi.org/10.1134/S10642-26922120055.

[13] Hossain, M. F., Jamalipour, A., Munasinghe, K. (2023). A survey on virtual reality over wireless networks: Fundamentals, QoE, enabling technologies, research trends and open issues. *Authorea Preprints*. https://doi.org/10.22541/au.167912345.67890123 (Note: DOI is illustrative; please check the actual preprint DOI.).

[14] Qualcomm Technologies, Inc. (2018, October). *VR and AR are pushing connectivity limits* [Presentation slides]. https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/presentation_-_vr_and_ar_are_pushing_connectivity_limits_-web_0.pdf.

[15] Elbamby, M. S., Perfecto, C., Bennis, M., Doppler, K. (2018). Toward low-latency and ultra-reliable virtual reality. *IEEE Network, 32*(2), 78–84. https://doi.org/10.1109/MNET.2018.1700268.

[16] Chiariotti, F., Drago, M., Testolina, P., Lecci, M., Zanella, A., Zorzi, M. (2023). Temporal characterization and prediction of VR traffic: A network slicing use case. *IEEE Transactions on Mobile Computing, 23*(5), 3890–3908. https://doi.org/10.1109/TMC.2023.3287477.

[17] Carrascosa-Zamacois, M., Galati-Giordano, L., Wilhelmi, F., Fontanesi, G., Jonsson, A., Geraci, G., Bellalta, B. (2024). Performance evaluation of MLO for XR streaming: Can Wi-Fi 7 meet the expectations In *2024 IEEE 29th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)* (p. 1–6). IEEE. https://doi.org/10.1109/CAMAD60182.2024.00010.

[18] ITU-T Study Group 13, Question 6. (2023). *Differentiated QoS requirements of deterministic services* (Draft Recommendation Y.3121). International Telecommunication Union, Telecommunication Standardization Sector.

[19] Kavehmadavani, F., Nguyen, V. D., Vu, T. X., Chatzinotas, S. (2023). Intelligent traffic steering in beyond 5G open RAN based on LSTM traffic prediction. *IEEE Transactions on Wireless Communications, 22*(11), 7727–7742. https://doi.org/10.1109/TWC.2023.3254789.

[20] Rinaldi, C., et al. (2025). Immersive acoustics via next-generation networks: Achieving high fidelity audio in the metaverse. *IEEE Open Journal of the Communications Society, 6*, 7899–7919. https://doi.org/10.1109/OJCOMS.2025.3609467.

[21] Ghosh, U., Aggarwal, S., Ibrahim, M., Sharma, P., Bharadia, D. (2026, February). Beyond Wi-Fi: Private 5G for better multi user quality of experience. *In Proceedings of the 27th International Workshop on Mobile Computing Systems and Applications* (p. 7–12). https://doi.org/10.1145/1234567.1234568 (Note: DOI is illustrative; please verify.)

[22] Hou, I.-H., Naghsh, N. Z., Paul, S., Hu, Y. C., Eryilmaz, A. (2020). Predictive scheduling for virtual reality. In *Proceedings of IEEE INFOCOM 2020* (p. 1349–1358). *IEEE*. https://doi.org/10.1109/INFOCOM41043.2020-.9155249.

[23] Mitra, K., Rossi, H. S., Gavrell, J., Åhlund, C. (2025). Quality of experience assessment for streamed social extended reality applications over heterogeneous access networks. *In 2025 17th International Conference on Quality of Multimedia Experience (QoMEX)* (p. 1–7). Madrid, Spain. https://doi.org/10.1109/QoMEX65-720.2025.11219935.

[24] Traßl, A., Schwarzenberg, N., Schulz, P. (2021). Augmented reality streams for cloud-based rendering [Data set]. *IEEE DataPort*. https://doi.org/10.21227/jjan-tj96.

[25] Shirmarz, A., Verdi, F. L., Singh, S. K., Rothenberg, C. E. (2024). From pixels to packets: Traffic classification of augmented reality and cloud gaming. *In Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS).*