Standardizing Multimedia QoE Telemetry from Telecommunications Networks for Open Analytics

Yifan Wang
University of New South Wales
Sydney, NSW, Australia
wangyifan.frank@unsw.edu.au

Minzhao Lyu
University of New South Wales
Sydney, NSW, Australia
minzhao.lyu@unsw.edu.au

Vijay Sivaraman
University of New South Wales
Sydney, NSW, Australia
vijay@unsw.edu.au

ABSTRACT

Telecommunications network operators routinely use active and passive network monitoring tools for observability into their networks, so they can ensure good quality of experience (QoE) for their users on real-time multimedia applications such as video streaming and online gaming. Commercial tools for network observability have typically been standalone appliances that export telemetry in proprietary formats, which limits network operators from coupling the data with emerging cloud-based analytics and AI platforms to unlock more value in areas such as network operations, customer care, and personalized product creation. In this paper, we aim to fill this gap by proposing a standardized schema for representation and export of QoE for a wide range of multimedia applications that can help bridge network observability tools with data analytics platforms. We design and implement an open-source tool by extending the widely used IPFIX and OpenTelemetry standards, and demonstrate a live implementation in a University campus network carrying thousands of application streams.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Network monitoring**;
• **Information systems** → **Data layout**;

KEYWORDS

Multimedia application, Quality-of-Experience, network telemetry

ACM Reference Format:
Yifan Wang, Minzhao Lyu, and Vijay Sivaraman. 2024. Standardizing Multimedia QoE Telemetry from Telecommunications Networks for Open Analytics. In *SIGCOMM Workshop on Emerging Multimedia Systems (EMS '24), August 4–8, 2024, Sydney, NSW, Australia.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3672196.3673400

1 INTRODUCTION

Telecommunications network operators seek visibility into their networks to understand how their customers are served. Maintaining high Net Promoter Score (NPS) and Customer Satisfaction Score (CSAT) are always important, so they need to know whether their customers are getting a smooth experience on their video streaming, online gaming, video conferencing, and other real-time multimedia applications. Network operators therefore routinely deploy tools to measure not only quality-of-service (QoS) metrics (*e.g.,* throughput, latency, packet drop rate) in different parts of their network, but also user experience or QoE (in terms of video resolutions, game lag spikes, voice glitches, etc.) of popular multimedia applications used by their customers. Knowing only network QoS delivered by the network cannot well represent the actual user experience of multimedia applications that have different network demands. For example, video streaming requires high bandwidth at several tens of Mbps to deliver high-resolution videos to smart TVs; online gaming imposes strict requirement on network latency and packet drop rate; video conferencing expects stably high bandwidth and low latency for real-time conversations.

To serve this important business need of telecommunications operators, as will be discussed in §2, vendors such as Cisco, Sandvine, Allot and Canopus offer network monitoring and visibility tools. As reported by [26], network visibility has become a $20B industry in 2022 and has an anticipated growth to $30B by 2027. In addition to active measurement methods that probe the network with sample sessions at certain test locations at different times, operators are increasingly adopting passive techniques that analyze traffic streams to get true measures and more comprehensive coverage of user experience.

Multimedia QoE data from the telemetry systems is exported to analytics teams for insights generation, often in the form of daily, weekly or monthly reports, which can be limited in timeliness and business value. Third-party data analytics platforms are emerging that specialize in extracting insights for business operations (*e.g.,* Reailize, selector.ai, Google's VertexAI) on cloud platforms by leveraging advances in predictive and generative AI. In order to plumb the QoE data into these platforms, it needs to be exported in standardized formats with metrics that are systematically defined in a commonly accepted data schema. Today's network observability tools either keep the data in databases with proprietary schemas or export it in non-standardized formats, limiting telecommunications operators from leveraging best-of-breed analytics capabilities emerging in the market.

There have been some prior research efforts [33] and industry initiatives [17] to define data schemas that foster open analytics of domain-specific data. For example, the MITRE corporation proposed structured data schema named CVE for cybersecurity vulnerability records, which has become a standard format for cybersecurity analytics [30]. In data networking, IPFIX (formerly known as NetFlow) and OpenTelemetry (OTel) have become commonly accepted telemetry formats for sharing structured data (defined

EMS '24, August 4–8, 2024, Sydney, NSW, Australia

ACM ISBN 979-8-4007-0711-7/24/08.
https://doi.org/10.1145/3672196.3673400

Vendor	Visibility	Granularity	Export format
Accedian [1]	App QoE	App session	Log file
Allot [2]	App QoE	Flow	Log file, database
Cambium [3]	Network QoS	Flow	Not provided
Canopus [4]	App QoE	App session	Log file, database
Sandvine [9]	App QoE	App session	Database, IPFIX
ThousandEyes [11]	Network QoS	Flow	Log file, OTel

Table 1: Six Representative vendors providing application network monitoring systems for ISPs.

in a systematic schema) from the measurement tool to the data analytics platform. However, to the best of our knowledge, they have not been extended to include application QoE data, which is the focus of this paper.

In this context, this paper designs a systematic schema defining multimedia application QoE data that can be exported by network observability systems to data analytics platforms leveraging existing standards such as IPFIX and OTel. Our specific contributions are three-fold:

Our **first contribution** (§3) develops a data schema that systematically categorizes telemetry metrics of multimedia application sessions that can be generated by network visibility systems. Our schema is designed to have three data blocks consisting of metadata, network statistics, and application-specific QoE. The app-specific QoE data block currently contains eight popular multimedia applications: on-demand video streaming, live video streaming, online gaming, voice/video calling, music streaming, cloud gaming, AR/VR, and social networking; it can be further extended easily for more applications.

Our **second contribution** (§4) formally defines our data schema via JSON and YAML files that have been reviewed by industry experts and are made publicly available at [34]. We also design and implement an open-source tool to export telemetry (from network observability systems) aligned with our data schema as standard IPFIX and OTel messages, which can be readily received and analyzed by cloud-based data platforms.

In the **third contribution** (§5), we integrate our data schema and telemetry export tool with a commercial network visibility system. We demonstrate its functionalities using ground-truth video streaming and gaming sessions in a University campus network, showcasing the process that converts non-standard telemetry to our data schema and exports them as IPFIX and OTel messages.

2 BACKGROUND AND RELATED WORK

***Network Monitoring and Visibility Systems*:** Network monitoring and visibility systems provide telecommunications operators with real-time visibility into user experience (QoE) of multimedia applications, such as on-demand video streaming [15, 16], live streaming [19, 24], online gaming [25], video conferencing [23, 27, 31], cloud gaming [12, 13, 21], and AR/VR [22]. These multimedia apps are increasingly important for network operators as they have a high impact on customer satisfaction and demand high network capabilities, such as high bandwidth and low latency. In Table 1, we list six commercial systems with their provided visibility types (either application QoE or network QoS), telemetry granularity (at the level of user application session or only network flow), and data export formats. Two systems (*i.e.,* Cambium and ThousandEyes) only provides QoS data of their monitored networks such as latency and flow throughput, while the other four vendors generate application-specific QoE metrics per app session from the monitored networks. For example, Canopus exports graphic resolution and frame rate of video streaming sessions and Sandvine measures game genres and gameplay lags for online gaming sessions. The generated measurement data from different systems are exported to network operators in diversified formats, mostly non-standardized/proprietary (*e.g.,* log file or entries in private time-series database).

***A Challenge in Multimedia Network Telemetry Analytics for the Data Community*:** While network visibility systems can provide valuable telemetry of multimedia sessions served by the monitored networks, the current use of non-standardized data formats, as discussed earlier, introduces analytical barriers for network operators to fully benefit from the telemetry. In the current industrial practice, either a well-trained internal data analytics team of the network operators or customer support team from the visibility vendors analyze the telemetry records. Other teams of data experts or relevant stakeholders can hardly leverage the valuable data due to lack of support and skills in processing the non-standardized/proprietary telemetry formats from various vendors. There are companies [7, 10] that offer strong expertise in extracting operational and business insight from data in standardized formats (*e.g.,* IPFIX and OpenTelemetry), which are either stored in local databases or cloud platforms such as Google GCP, Amazon AWS. There are also data specialists on cloud data platforms like Databricks [6] who are available to perform data analytics tasks for monetary rewards. While the business and operational benefits for network operators to leverage the potentially massive data analytics skills and capability on those platforms are huge [26], having non-standardized/proprietary telemetry data formats from network monitoring systems, which are not commonly supported by the data community, makes the objective less feasible in practice.

***Standard Telemetry Export Format*:** Commercial companies and data specialist communities, such as Cisco and CNCF [5], have been working on network data in standardized formats such as IPFIX [18] and OpenTelemetry [8] (OTel). For example, telemetry of network flows such as flow metadata and volumetric attributes that are generated by network visibility systems can be encapsulated as information elements in IPFIX messages. These standard IPFIX messages are exported to one or multiple receivers (*e.g.,* data experts and companies) specializing in data analytics. Note that both IPFIX and OTel require both senders (*i.e.,* telemetry system) and receivers to agree on the information they would like to exchange with a predefined template (*i.e.,* data schema). Therefore, in the next section, we develop a systematic data schema to comprehensively capture metrics related to multimedia application user experience that can be exported by network telemetry systems.

3 TELEMETRY DATA SCHEMA AND EXPORTING MECHANISM

In this section, we first (in §3.1) design a process of exporting QoE telemetry for multimedia applications from being generated

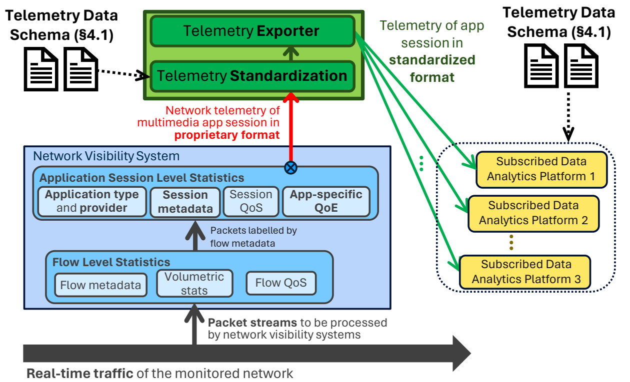

Figure 1: Process of exporting QoE telemetry from network visibility systems for open analytics.

Application type	Popular provider	Example QoE metric
On-demand video streaming	YouTube, Netflix, Hulu	Resolution, frame rate
Live video streaming	Twitch, YouTube Live	Startup delay, video stall
Online gaming	Call of Duty, Fortnite	Gameplay lag, game freeze
Voice/video calling	WhatsApp, Zoom, Teams	Voice mute, frame drop
Music streaming	Spotify, Apple Music	Audio bitrate, loading delay
Cloud gaming	GeForce NOW, xCloud	Resolution, frame rate
AR/VR	VRChat, RecRoom	Activity lag, motion freeze
Social networking	X, Facebook, Instagram	Page load time

Table 2: Eight Types of multimedia applications covered by our data schema.

by network visibility systems to being received by data analytics platforms. We then (in §3.2) design a data schema that standardizes QoE metrics for eight types of multimedia applications as listed in Table 2. For simplicity, we discuss video streaming and online gaming in details.

3.1 Network Telemetry Exporting Processes

The telemetry generation and exporting process involves three components as color-coded by blue, green, and yellow blocks in Fig. 1 for visibility system, telemetry standardization and exporting modules, and analytics platforms, respectively.

As shown by the blue block, commercial **network visibility systems** process real-time traffic of the monitored networks and generate flow-level network QoS statistics such as packet rate and latency, which can be further aggregated to application session-level telemetry (*e.g.,* video resolution and gameplay lag) depending on the specific system capabilities. In this paper, we are particularly interested in session-level telemetry of multimedia applications, which is generated when a new app session is detected, periodically, when a special event occurs such as a sudden drop or rise of an important metric, or when a session is finished.

To make the generated QoE telemetry of multimedia application sessions available and readily consumed for open analytics, we place a **telemetry standardization module** (the bottom green block in Fig. 1) that subscribes to non-standardized/proprietary telemetry by reading log files or retrieving entries from databases of visibility systems. This module coverts the telemetry by referring

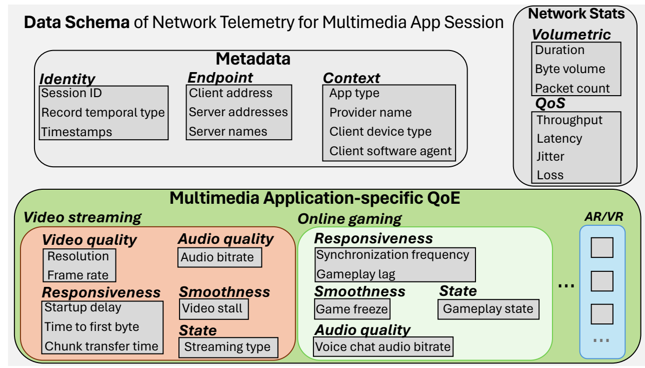

Figure 2: Our data schema for network telemetry of multimedia applications. The technical JSON and YAML files have been reviewed by industry experts and are publicly available at [34].

to a systematic data schema (to be detailed in §3.2), which are then encapsulated as standard telemetry export formats (*e.g.,* IPFIX and OTel as implemented in this paper) by the **telemetry exporter** represented by the top green box to **data analytics teams** who are authorized to subscribe.

3.2 Data Schema of Network Telemetry for Multimedia Applications

Our data schema that standardizes QoE telemetry of multimedia application sessions from network visibility systems contains three blocks, including metadata, network statistics, and application-specific QoE metrics as visually depicted in Fig. 2. Our data schema is documented using JSON and YAML files as will be discussed in §4.

3.2.1 Metadata. Metrics in the metadata block do not contain network QoS or application QoE metrics, but provide labels for telemetry grouping and analytics. First of all, **identity** information of each multimedia application session contains a unique identifier generated for the particular app session (*i.e.,* session ID), the record temporal type showing whether the telemetry record is generated at the start of the session, periodic points, event-triggered, or the end of the session, and its timestamps. The **time granularity** of telemetries is flexibly defined by the connected network visibility systems, such as based on static or dynamic interval (*e.g.,* every one minute) or specific events (*e.g.,* the start, end, or certain state changes) of a monitored session. IP addresses of the client and server(s) and server name(s) of an app session are defined in the **endpoint** subcategory. We note that a multimedia app session can have multiple servers and server names. For example, a video streaming session can receive media data from multiple cache servers, and an online gaming session can be hosted by different servers for game matching and gameplay. Therefore, server IP addresses and names are captured as array data type in our data schema. In the **context** subcategory, application type (*e.g.,* video streaming or online gaming), provider name (*e.g.,* YouTube or Call of Duty), client device type (*e.g.,* macOS or iOS), and software agent (*e.g.,* Chrome browser or native app)

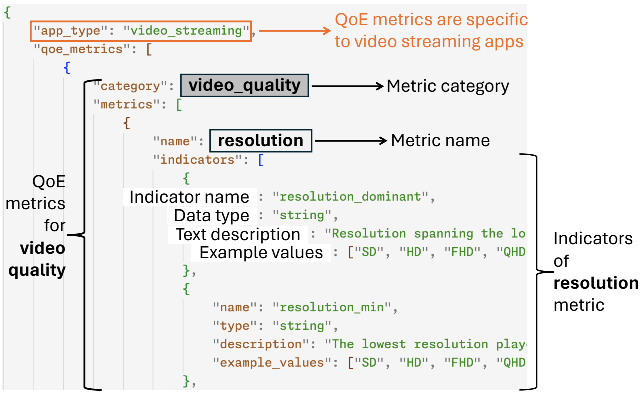

Figure 3: An example snippet of JSON files describing the metrics of video resolution.

are defined to provide contextual labels for the measured QoS and QoE data.

3.2.2 Network Statistics. This category describes an application session by its aggregated volumetric usage and instantaneous network quality-of-service (QoS) metrics. The aggregated volumetric usage contains session duration, total volume and packet count. QoS metrics in our schema including throughput, latency, jitter, and packet loss have been well studied and used by the Internet community. We note that an app session can be supported by multiple concurrent network flows. Therefore, the value of each QoS metric can be represented by the summation of all flow-level values (*e.g.,* for throughput) or other statistical indicators (*e.g.,* median) as implemented by the specific network visibility system.

3.2.3 Application-specific QoE. This category of metrics contains QoE telemetry that can be measured from network traffic for each multimedia application type. Here we discuss in details the metrics for video streaming and online gaming each as a subcategory. Metrics for the other application types in Table 2 are omitted in this paper for simplicity and available at [34].

For **video streaming**, our schema contains all metrics for video quality, audio quality, responsiveness, smoothness, and state that can be measured by commercial network visibility systems or reported by related research works [15, 24, 32]. Specifically, video quality (resolution and frame rate) and audio quality (audio bitrate) indicate the clearness and clarity of the instantaneous media playback. Responsiveness (*i.e.,* startup delay, time-to-first-byte and chunk transfer time) measures the amount of temporal delays an user may experience during a video session. Smoothness (*i.e.,* video stall) shows how frequent a video is stuck during a measured period of time. The state (video streaming type) can be either Video-on-Demand (VoD) or live streaming, which can interchange during one video session, *e.g.,* when a user watching a live stream drags the video progress bar backwards causing the state to be changed to VoD.

The subcategory of **online gaming** covers metrics for audio quality of in-game voice chat, responsiveness (*i.e.,* gameplay lag and synchronization frequency), smoothness of user activity (*i.e.,* occurrence of motion freeze), and state (*e.g.,* intensive gameplay or less-intense phases between rounds of the game) that are determined by or can be measured from network traffic according to prior works [25, 28].

4 FORMALLY DEFINING DATA SCHEMA IN STANDARD FORMATS

We now discuss how the data schema is formally defined as standard JSON and YAML files that are in wide use by the industry bodies (§4.1) and how the respective real-time multimedia telemetry is exported as standard IPFIX and OTel messages (§4.2).

4.1 Defining Data Schema as JSON and YAML Files

To make our QoE data schema usable by the industry bodies, we formally describe it in two standard formats, *i.e.,* JSON and YAML, that have been popular in defining industry API protocols and information frameworks [20, 29]. Both formats can well preserve the hierarchical grouping of metrics in our data schema. In addition to the name of each metric as shown in Fig. 2, we specify their data type (*e.g.,* int, string or array), provide description, and suggest example telemetry values as reference. Note that QoS and QoE metrics, *e.g.,* graphic resolution, can be represented by various statistical indicators such as the dominant level during the reporting period, the lowest resolution level, and the duration for each level of resolution. Therefore, as shown in the example snippet Fig. 3 for video graphic resolution, in our definition, each QoS and QoE metric is measured by a list of indicators that are commonly used by current commercial telemetry systems.

4.2 Exporting Telemetry as IPFIX and OTel Messages

We develop an open-source tool (*i.e.,* telemetry exporter module in Fig. 1) to export real-time telemetry (after being standardized by our JSON/YAML data schema) as two commonly used data export formats, namely IPFIX and OTel [34]. The exported messages are then received by standard IPFIX/OTel receivers of data analytics teams for business and operational insights.

***IPFIX*:** IPFIX has been commonly used to export network telemetry data by the industry especially for flow statistics and QoS. An IPFIX message contains independent data entries, called information elements (IEs) as in a plain data table. Each IE contains the name, data type, and semantic information of the metric to be exported. Therefore, after receiving a telemetry that follows our standard data schema in either JSON or YAML format, the exporting module will convert each metric as a plain table entry without hierarchical grouping (one example snippet is shown in Fig. 3), which are then sent to the subscribed analytics platform as IPFIX messages. As shown in Fig. 1, IPFIX receivers in the analytics platform can restore the hierarchical relationships defined in the data schema. In addition to data messages containing specific metric values, IPFIX also allows an exporter to specify which metrics will be exported to the receivers via "template" messages. Therefore, prior to the transmission of telemetry data messages for a newly

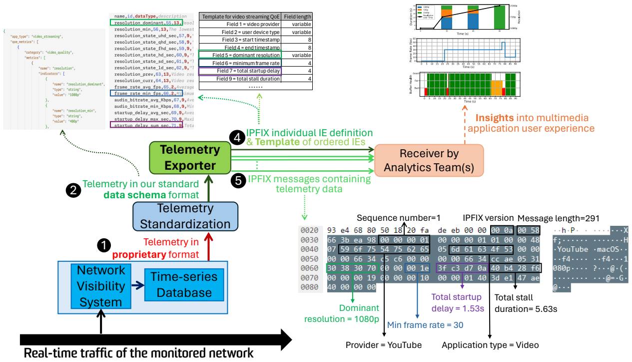

Figure 4: Integrating our telemetry standardization and exporter tool with a commercial network observability system.

detected multimedia app session, our exporting tool will send "template" messages to inform the receivers of which metrics defined in the schema will be sent for the particular session.

***OTel*:** OTel is a format for exporting performance information of cloud-native applications. OTel establishes dedicated data streams for each defined metric between the sender and subscribed receivers over an entire monitoring session, *e.g.,* from the start to end of a multimedia app session. Therefore, unlike our exporting implementation for IPFIX that sends all metrics in one data message, we create concurrent data pipes for telemetry metrics. Dynamic QoS and QoE metrics are assigned as OTel's "metric" type which can be updated over time, timing metrics in our data schema are exported as OTel's "span" data type, and static metadata fields are specified as OTel's "attribute" type.

5 DEMONSTRATION WITH A COMMERCIAL OBSERVABILITY SYSTEM

We have integrated our open-source telemetry standardization and exporter tool with a commercial network visibility system that is used by Telecommunications for multimedia application customer insights, as visually shown in Fig. 4. The visibility system streams its measured telemetry into a proprietary data pipeline (step ①). Our tool subscribes to the streamed telemetry (step ③) that is formatted by our data schema in step ②. The telemetry is then exported as IPFIX messages to the standard IPFIX receiver of data analytics teams (step ④,⑤). We deployed our system in a University campus network and evaluated it on many application types and flows. In what follows, we use two ground-truth sessions played in our network and measured by the visibility system to showcase how the QoE telemetry of the two sessions is standardized and exported by our tool, and analyzed by data specialists.

***A YouTube Video Session*:** In this video session, one of the authors watched a YouTube video for 90 seconds on a macOS laptop via Chrome browser. As shown in Fig. 5(a), the video user experienced various types of quality drops such as long startup delay for three seconds, video stalls for a few seconds, resolution dropping from 4K to 1080p and 480p, and changed frame rates. The commercial network visibility system is capable of reporting these session-level metrics (that are stored in its time-series database

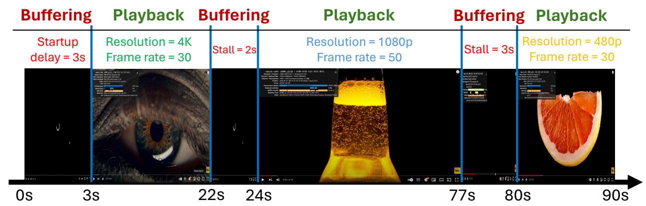

(a) A YouTube video session on the user device.

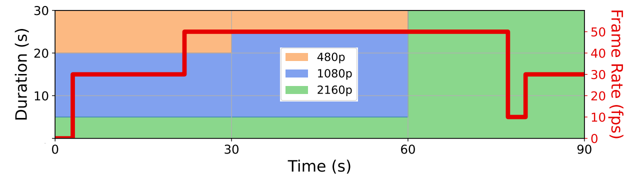

(b) Video resolution and frame rate.

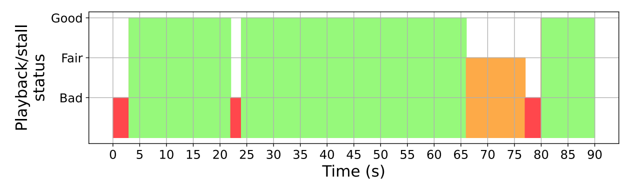

(c) Video playback/stall status.

Figure 5: Exporting telemetry for (a) a YouTube session on the user device as (b), (c) standard QoE metrics received by the analytics team.

in proprietary formats) at the frequency of thirty seconds. Our telemetry standardization and exporting tool continuously extracts new telemetry entries, converts them with our data schema as defined in JSON files, which are then exported as IPFIX messages. A screenshot of an IPFIX message snippet captured using Wireshark is provided as step ④ in Fig. 4. Telemetry encapsulated as IPFIX messages are received by data specialists who will perform experience insight analytics. As a demonstrative example, we visualize the received values of three QoE metrics, *i.e.,* resolution and frame rate in Fig. 5(b) and video playback/stall status in Fig. 5(c), for video streaming application. The stacked color blocks show the total duration of video playback at the respective resolutions (*i.e.,* 480p, 1080p, and 2160p) for each 30-second interval from 0-30s, 30-60s, and 60-90s. The red line in Fig. 5(b) illustrates the reported video frame rate over the entire session. The video playback/stall status shown in Fig. 5(c), as reported by periodic telemetries and those triggered by state change events (*e.g.,* buffering, resumption of playback), is described in three levels of bad, fair, and good.

***A Call-of-Duty Online Gaming Session*:** To demonstrate our telemetry export for online gaming application, one of the authors played an online multiplayer shooting game (Call of Duty: Modern Warfare) that requires low gameplay lag and small jitters for good user experience. In the game session, we started by entering the lobby for matchmaking, and then had both smooth and choppy (*i.e.,* not smooth) latency experience as labeled in Fig. 6(a). For online gaming sessions, the network visibility system sends telemetry

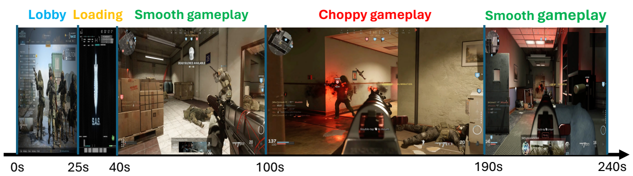

(a) A CoD gameplay session on the user device.

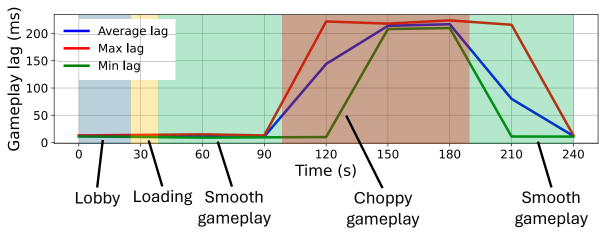

(b) Gameplay lags and game states received by the analytics team.

Figure 6: Exporting telemetry for (a) a Call of Duty online gaming session on the user device as (b) standard QoE metrics received by the analytics team.

periodically per thirty seconds and also when game state changes (*e.g.,* from matchmaking to gameplay or from smooth gameplay experience to choppy experience). After receiving IPFIX messages that carries QoE metrics formatted by our data schema, the data specialists can provide insights into user experience of this particular game session, such as via a visual graph in Fig. 6(b). The figure shows average, max, and min gameplay lags per reporting interval and their game states.

6 DISCUSSION

The objective of our developed data schema and convertating tool is to provide a unified framework describing multimedia QoE and QoS metrics being measured and reported by network visibility systems. In this sense, new metrics and multimedia application types being introduced to the domain will require an expansion or revision of particular modularized block(s) in our data schema.

We also acknowledge that there are future aspects to systematically evaluate the performance, scalability, and data security of our tool. Specifically, performance of our tool such as the timing overhead and computing resource consumption and their potential spaces for optimization can be studied with rigorous analysis. The scalability of our tool can be evaluated and further enhanced with state-of-the-arts system design principles including splitting of control-plane architecture and minimizing client-side sensing as highlighted in [14]. In addition, our current deployment is built on a single server covering telemetries from a campus network. Deployment at a commercial grade ISP will require collaborations of multiple processing nodes (servers).

Last, our efforts for opening the analytics of multimedia application telemetries for third-party and cloud data experts will inevitably introduce concerns on data security. To what extend the data can be shared and analyzed for maximized business value without compromising privacy of multimedia users require discussions. Also, data protection and access control against unauthorized subscribers especially on cloud platforms are important considerations.

7 CONCLUSION

We develop a systematic data schema that standardizes multimedia QoE data from commercial network visibility systems. Telemetry aligned with the data schema is exported as IPFIX or OpenTelemetry messages to data specialists for telecommunications network operational and business insights. We demonstrate a full integration of our data schema with a commercial network visibility system, showcasing standardized QoE telemetry data of video streaming and online gaming sessions being exported for open analytics.

ETHICS

We have obtained ethics clearance (UNSW Human Research Ethics Advisory Panel approval number HC211007) to analyze campus network traffic for application usage behaviors without identifying user identities. No attempt is made to extract or reveal any personal user information.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful feedback and valuable suggestion on future work. This work is supported by the Australian Government's National Industry PhD Program award reference number 35063 and Cooperative Research Centres Projects (CRC-P) Grant CRCPXIV000099.

REFERENCES

[1] 2024. Accedian. https://accedian.com/. Accessed: 2024-04-02.
[2] 2024. Allot LTD - Network Inteligence & Security Solutions and Protection. https://www.allot.com/. Accessed: 2024-04-02.
[3] 2024. Cambium Networks | Connectivity That Just Works. https://www.cambiumnetworks.com/. Accessed: 2024-04-02.
[4] 2024. Canopus: AI for Telcos. https://canopusnetworks.com/. Accessed: 2024-04-02.
[5] 2024. Cloud Native Computing Foundation. https://www.cncf.io/. Accessed: 2024-05-03.
[6] 2024. Data Streaming | Databricks. https://www.databricks.com/product/data-streaming. Accessed: 2024-05-02.
[7] 2024. Digital Media and Entertainment Monitoring | Datadog. https://bit.ly/datadog-telecoms. Accessed: 2024-04-30.
[8] 2024. OpenTelemetry. https://opentelemetry.io/. Accessed: 2024-05-01.
[9] 2024. Sandvine: The App QoE Company. https://www.sandvine.com/. Accessed: 2024-04-02.
[10] 2024. Telecom Data Cloud | Snowflake for Telecommunications. https://bit.ly/snowflake-telecoms. Accessed: 2024-04-02.
[11] 2024. ThousandEyes: Digital Experience Monitoring. https://www.thousandeyes.com/. Accessed: 2024-04-02.
[12] Carlos Baena, O.S. Peñaherrera-Pulla, Raquel Barco, and Sergio Fortes. 2023. Measuring and Estimating Key Quality Indicators in Cloud Gaming Services. *Computer Networks* 231 (2023), 109808.
[13] Sandeepa Bhuyan, Shulin Zhao, Ziyu Ying, Mahmut T. Kandemir, and Chita R. Das. 2022. End-to-end Characterization of Game Streaming Applications on Mobile Platforms. *Proc. ACM Meas. Anal. Comput. Syst.* 6, 1, Article 10 (feb 2022), 25 pages.
[14] Aditya Ganjam, Junchen Jiang, Xi Liu, Vyas Sekar, Faisal Siddiqi, Ion Stoica, Jibin Zhan, and Hui Zhang. 2015. C3: Internet-Scale Control Plane for Video Quality Optimization. In *Proc. USENIX NSDI*. Oakland, CA, 131–144.
[15] Craig Gutterman, Katherine Guo, Sarthak Arora, Xiaoyang Wang, Les Wu, Ethan Katz-Bassett, and Gil Zussman. 2019. Requet: Real-time QoE Detection for Encrypted YouTube Traffic. In *Proc. ACM MMSys*. Amherst Massachusetts, 48–59.
[16] Hassan Habibi Gharakheili et al. 2019. iTeleScope: Softwarized Network Middle-Box for Real-Time Video Telemetry and Classification. *IEEE Transactions on Network and Service Management* 16, 3 (Sept. 2019), 1071–1085.

[17] International Telecommunication Union. 2024. ITU Telecommunication Standardization Sector. https://www.itu.int/en/ITU-T/Pages/default.aspx. Accessed: 2024-05-03.

[18] Internet Assigned Numbers Authority. 2024. IP Flow Information Export (IPFIX) Entities. https://www.iana.org/assignments/ipfix/ipfix.xhtml. Accessed: 2024-04-04.

[19] Sina Keshvadi and Carey Williamson. 2021. An Empirical Measurement Study of Free Live Streaming Services. In *Proc. PAM*, Vol. 12671. Virtual Event, 111–127.

[20] Eliot Lear, Ralph Droms, and Dan Romascanu. 2019. *Manufacturer Usage Description Specification.* RFC 8520. doi:10.17487/RFC8520.

[21] Minzhao Lyu et al. 2024. Network Anatomy and Real-Time Measurement of Nvidia GeForce NOW Cloud Gaming. In *Proc. PAM.* Virtual Event.

[22] Minzhao Lyu, Rahul Dev Tripathi, and Vijay Sivaraman. 2023. MetaVRadar: Measuring Metaverse Virtual Reality Network Activity. *Proc. ACM Meas. Anal. Comput. Syst.* 7, 3 (Dec. 2023).

[23] Kyle MacMillan, Tarun Mangla, James Saxon, and Nick Feamster. 2021. Measuring the Performance and Network Utilization of Popular Video Conferencing Applications. In *Proc. ACM IMC.* Virtual Event, 229–244.

[24] Sharat Chandra Madanapalli et al. 2021. ReCLive: Real-Time Classification and QoE Inference of Live Video Streaming Services. In *Proc. IEEE/ACM IWQoS.* Tokyo, Japan, 1–7.

[25] Sharat C. Madanapalli et al. 2022. Know Thy Lag: In-Network Game Detection and Latency Measurement. In *Proc. PAM.* Virtual Event.

[26] MarketsandMarkets. 2022. Network Monitoring Market. https://bit.ly/network-monitoring-report. Accessed: 2024-05-03.

[27] Oliver Michel, Satadal Sengupta, Hyojoon Kim, Ravi Netravali, and Jennifer Rexford. 2022. Enabling Passive Measurement of Zoom Performance in Production Networks. In *Proc. ACM IMC.* Nice, France.

[28] Steven Schmidt, Saman Zadtootaghaj, Saeed Shafiee Sabet, and Sebastian Möller. 2021. Modeling and Understanding the Quality of Experience of Online Mobile Gaming Services. In *Proc. QoMEX.*

[29] The Linux Foundation. 2022. OpenAPI Initiative. https://www.openapis.org/. Accessed: 2024-05-04.

[30] The MITRE Corporation. 2024. Glossary | CVE. https://www.cve.org/ResourcesSupport/Glossary. Accessed: 2024-05-03.

[31] Matteo Varvello, Hyunseok Chang, and Yasir Zaki. 2022. Performance Characterization of Videoconferencing in the Wild. In *Proc. ACM IMC.* Nice, France, 261–273.

[32] Sarah Wassermann, Michael Seufert, Pedro Casas, Li Gang, and Kuang Li. 2020. ViCrypt to the Rescue: Real-Time, Machine-Learning-Driven Video-QoE Monitoring for Encrypted Streaming Traffic. *IEEE Transactions on Network and Service Management* 17, 4 (Dec. 2020), 2007–2023.

[33] Sascha Welten, Laurenz Neumann, Yeliz Ucer Yediel, Luiz Olavo Bonino da Silva Santos, Stefan Decker, and Oya Beyan. 2021. DAMS: A Distributed Analytics Metadata Schema. *Data Intelligence* (2021).

[34] WYF99. 2024. QoE-telemetry-standardized. https://github.com/WYF99/QoE-telemetry-standardized.