

Feasibility of Immersive Concert Streaming over a Private 5G Network

Edoardo Pilia¹, Gianluca Fadda¹, Maurizio Murrone^{1,2}, Nicușor Amarie³

¹DIEE, University of Cagliari, 09123 Cagliari, Italy

²CNIT, University of Cagliari, 09123 Cagliari, Italy

³Transilvania University of Brasov, Brasov, Romania

{edoardo.pilia, gianluca.fadda, maurizio.murrone}@unica.it, {nicusor.amarie}@unitbv.ro

Abstract—The immersive live music experience represents one of the most resource-demanding applications for networked audio, combining high-bitrate media delivery with strict requirements on temporal stability and audio-visual synchronization. While the video dimension of immersive concerts has received growing attention, the delivery of high-channel-count spatial and multi-track audio remains comparatively under-characterized over real cellular infrastructure.

This work investigates the feasibility of an immersive music concert experience over a private 5G Standalone (SA) network, in the context of the HEAT (Hybrid Extended reAliTy) EU Horizon Europe project. In this work, we drive a private Open Radio Access Network testbed with a multi-track concert signal encoded as independent Opus audio streams and delivered end-to-end through WebRTC. We characterize the system across an audio channel-count sweep from 4 to 64 simultaneous tracks, and we validate the results on two independent instrumentation layers: application (WebRTC metrics) and the RAN (E2SM-KPM counters exposed by a near-real-time RAN Intelligent Controller xApp). Results show that the offered load scales linearly with the number of channels, maintaining stable round-trip-time, jitter, and packet loss.

The findings demonstrate that a private 5G network transparently sustains immersive multi-channel concert audio, and quantify the structural overhead of a separate-track delivery approach.

Index Terms—Internet of Sounds, Musical Metaverse, 5G, Open RAN, Immersive Audio

I. INTRODUCTION

Immersive live music is emerging as one of the most demanding applications at the confluence of Extended Reality (XR) and networking [1], [2]. A faithful concert experience depends not only on immersive video, but especially on the delivery of high-fidelity, spatially organized audio. Unlike broadcast music, this audio is frequently carried as many individual instrument stems rather than a single muxed track, so that each source can be spatialized, balanced or manipulated independently at the receiver. This multichannel structure places concert audio within the Internet of Sounds and turns its delivery into a networking problem in its own right: not one stream, but tens of concurrent streams that must be sustained together [3], [4].

This study was carried out within the “HEAT – Hybrid Extended reAliTy” Project GA 101135637 funded by the EU Horizon Europe Framework Programme (HORIZON). This manuscript reflects only the authors’ views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

Delivering such content at a live event requires low-latency, high-throughput connectivity that conventional Wi-Fi may struggle to guarantee in multi-user scenarios due to contention [5], typical in crowded venues. Private 5G networks are an appealing alternative: they provide managed scheduling, Quality-of-Service (QoS) support and edge proximity, and 3GPP has explicitly studied XR traffic and its stringent rate and latency requirements over 5G [6], [7]. Programmable Open Radio Access Network (O-RAN) deployments expose fine-grained radio telemetry and control [8].

Most of this work, however, characterizes the network with immersive video and point-cloud workloads, or with synthetic traffic generators, rather than with a real multichannel audio application [5]. A synthetic bulk flow and tens of low-bitrate, independently paced audio streams place very different demands on the radio scheduler and on the receiver, and the latter is precisely the regime of interest for immersive concert audio. How a multichannel audio application behaves over a 5G SA network—and how it scales as the number of audio channels grows—therefore remains largely uncharacterized.

This paper addresses that gap. We deliver a real multi-track concert signal over a private 5G O-RAN testbed and characterize its feasibility across an increasing number of radio channels. We find that the offered load scales linearly with the channel count, with end-to-end latency and jitter remaining load-independent, and the RAN contributes sub-millisecond radio-link-control (RLC) delay throughout. Our specific contributions are:

- an end-to-end pipeline that streams real concert audio as independent per-channel Opus streams over WebRTC across a private 5G SA network;
- an experimental channel-count sweep (4, 8, 16, 32 and 64 simultaneous audio tracks, three repetitions each) driven by real application-layer content rather than synthetic traffic;
- a two-layer validation methodology that reconciles application/transport statistics with RAN-side E2SM-KPM counters, so that no single measurement layer is relied upon in isolation;
- a quantification of the structural per-stream protocol overhead of the separate-track delivery approach.

The remainder of this paper is organized as follows. Sec-

tion II reviews related work on network immersive audio and on immersive media over 5G. Section III describes the testbed, the audio pipeline and the instrumentation. Section IV reports the channel-count sweep and the cross-layer validation. Section V concludes and outlines future work.

II. RELATED WORKS

Real-time audio over networks has a long history in the context of Networked Music Performance (NMP), where geographically separated musicians interact under a strict sub-30ms one-way latency [9]. Systems such as JackTrip established the reference approach of sending low-latency, uncompressed audio over UDP with small jitter buffers [10], later extended to wireless and embedded settings [11] and, more recently, to browser-based delivery via WebRTC [12]. These systems are fundamentally bidirectional and peer-to-peer: their defining constraint is the closed-loop latency between performers. Recent work has begun to address NMP over 5G access specifically [13]. The scenario we address is different. We consider the downstream distribution of an immersive concert to an audience—a one-to-many problem in which the dominant requirement is sustaining many concurrent audio streams with stable delivery, rather than closing an interactive loop. Minimizing latency remains desirable, and low-latency network is a broad concern beyond music [14], but the scaling behavior with the number of channels becomes central.

That scaling is precisely what immersive spatial audio makes acute. In Ambisonics-based delivery, increasing the spatial order improves immersion but multiplies the number of audio channels, escalating both bandwidth and susceptibility to network impairments; adaptive schemes have been proposed to trade spatial order against available throughput [15]. Opus, mandatory to implement in WebRTC, is a common choice for compressing such multichannel spatial audio [16]. This channel count versus network tradeoff is the axis we characterize experimentally over a private 5G network.

A. Media delivery over O-RAN

The delivery of immersive media over 5G has been studied primarily for the visual modality. 3GPP has explicitly modeled XR traffic and its rate and latency requirements [6], [7], [17], and the broader challenges of XR-over-wireless have been surveyed in [18]. Concrete systems include split rendering with L4S for latency-critical interactive XR [19], cloud VR validated over 5G in industrial scenarios [20], and 5G-ready volumetric video streaming platforms [21]. On the infrastructure side, the O-RAN paradigm exposes disaggregated RAN components and fine-grained telemetry through the E2 interface [8]; open stacks such as OpenAirInterface (OAI) [22] together with the near-real-time RIC (Near-RT RIC) frameworks and their KPM service models [23] make RAN-internal measurements accessible to applications. We build on this capability to validate application-level results against RAN counters, but—unlike the above works, which target video or

volumetric content—our workload implements multichannel audio.

B. Immersive Media in the HEAT Project

Within the Hybrid Extended reAliTy (HEAT), several efforts target the immersive concert from complementary angles. The vision of interactive, distributed hybrid XR experiences is outlined in [24], while point-cloud streaming over 5G has been characterized in terms of Quality of Experience (QoE) [25] and made practical for a standalone headset through a wearable 5G-to-USB gateway [5]. On the audio side, an XR application that streams a musicians' multitrack mixer as spatial audio has been proposed [26]. The present work is the empirical complement to that application: whereas [26] proposed the immersive multichannel audio system, we characterize the feasibility and scaling of delivering such content over 5G networks, validated across application, transport and radio layers.

In summary, networked-audio research has focused on bidirectional NMP and immersive-media-over-5G research has focused on video and volumetric content. The behavior of a real, high-channel-count immersive audio application over a 5G network has not been characterized yet. The target of this paper is to fill this research gap.

III. SYSTEM IMPLEMENTATION

The workload characterized in this work originates from the HEAT concert pilots, in which a live musical performance is captured and distributed to a remote audience wearing XR devices. In this setting each instrument is carried as a separate stem, so the listener's renderer can place every source in the virtual space and adjust the balance according to the listener's position. Therefore, the network is required to sustain as many concurrent audio streams as the performance has sources.

Fig. 1 shows the end-to-end system, organized in three layers: a content/source layer, a network layer based on a private O-RAN 5G SA testbed, and an endpoint/destination layer. The concert audio is generated by the source server, injected into the network and delivered downlink across the 5G segment to the receiving host, reproducing the one-to-many distribution scenario of Section I. In the current scenario, both endpoints are host workstations: the source is attached on the network side of the 5G segment, while the receiving host itself acts as the 5G UE through a directly attached cellular module.

A. Radio Access Network

The private 5G SA network is a self-contained O-RAN deployment based on OAI [22]. The 5G Core and the gNB Central and Distributed Units (CUs, DUs) run on a dedicated workstation (Intel Core i9, 64 GB RAM, Ubuntu 22.04 LTS with real-time kernel), connected to a commercial VVDN Low Power Radio Unit (LPRU) over a 10 Gbps optical fronthaul using the O-RAN 7.2x split. The user equipment is a Quectel RM530N-GL 5G Module on an M.2-to-USB hat, operated in SA mode and managed over USB through the MBIM interface. Radio parameters are summarized in Table I. A near-real-time

TABLE I
O-RAN TESTBED CONFIGURATION.

Parameter	Value
Stack	OpenAirInterface
Core / gNB host	Intel Core i9, 64 GB RAM, Ubuntu 22.04 (RT kernel)
Radio Unit	VVDN Low-Power RU (LPRU)
Fronthaul	10 Gbps optical, O-RAN 7.2x split
Band / Duplex	n77, Time Division Duplex
Center frequency	4.0 GHz
Channel bandwidth	100 MHz (273 PRB)
Subcarrier spacing	30 kHz ($\mu = 1$), 0.5 ms slot
TDD pattern	5 ms periodicity: 7 DL + 2 UL slots, special slot 6 DL / 4 guard / 4 UL symbols
Antenna config.	2T2R, single layer
Synchronization	IEEE 1588v2 PTP
UE	Quectel RM530N-GL (MBIM), SA mode
RAN Monitoring	Near-RT RIC xApp, E2SM-KPM, 1 s period

RIC hosts an xApp that subscribes to E2SM-KPM measurements and exposes per-UE counters—throughput, RLC delay, assigned Physical Resource Blocks (PRBs)—at a one-second reporting period.

During the experiments both UEs were located approximately 5 m from the radio unit in line of sight.

B. Audio Pipeline

The concert audio is delivered by a Unity application built on `com.unity.webrtc`. Each audio channel is carried as an independent WebRTC audio track encoded with Opus at 64 kbps CBR and 20 ms framing, i.e. an approximately 160-byte payload at 50 packets per second per track (RTP payload type-96). The number of tracks placed on the wire is fixed at the connection setup and is not inferred from the media loaded at the receiver; changing it requires renegotiation. The channel count is therefore an explicit parameter equal to N independently encoded streams.

The source material is a single-instrument concert track from a previous HEAT experimentation (24-bit, 44.1 kHz, mono; peak -5.9 dBFS, crest factor 19.7 dB, resampled to 48 kHz before Opus encoding). To realize a controlled channel-count sweep, each condition transmits N instances of this stem, each independently encoded. Because the offered load of a separate-track pipeline is governed by per-stream encoding and packetization, this reproduces a representative offered load for N concurrent channels. The highest condition comprises 63 pre-recorded tracks plus one live-capture (ASIO) track, for 64 simultaneous RTP streams; the live-capture path was inactive during the runs, so the transmitted content is entirely pre-recorded multi-track audio.

All N tracks are carried within a single WebRTC Peer-Connection, so that they share one transport path and one congestion-control context, matching the deployment in which a listener subscribes to the full ensemble rather than to individual sources. A constant-bitrate encoder configuration was chosen in preference to variable bitrate so that the offered load is determined by the channel count, making the sweep interpretable as a controlled load parameter.

TABLE II
SUMMARY OF RESULTS ACROSS THE CHANNEL-COUNT SWEEP.

N	Application layer				RAN layer		
	Audio (kbps)	RTT (ms)	Jitter (ms)	Loss (avg %)	DL thp. (kbps)	UL thp. (kbps)	RLC (μ s)
4	259	34	5.0	0.004	612	415	190
8	516	34	6.2	0.076	1005	788	137
16	1029	34	5.9	0.026	1769	1537	246
32	2062	36	5.8	0.064	3301	3033	293
64	4109	34	8.2	0.164	6324	6016	449

This construction is deliberately conservative. Because Opus operates at a fixed 64 kbps CBR with 20 ms framing, the offered load of a separate-track pipeline is determined by per-stream encoding and packetization, and is independent of the musical content of each stem. What the use of identical copies does affect is the arrival pattern: encoding N identical signals in lockstep causes the sender to emit synchronized bursts of N packets every 20 ms, whereas decorrelated stems from a real ensemble would produce smoother, temporally spread arrivals. The measured burstiness is thus an upper bound on what a genuine musical performance would impose on the radio scheduler.

C. Measurement Instrumentation

Two independent instrumentation layers are used. At the application layer, WebRTC `getStats` at the Host side provides round trip time, jitter, packet loss, per-track and aggregated bitrate, and packet counts. At the radio layer, the E2SM-KPM xApp provides downlink and uplink throughput and RLC delay for both the Host and the Artist UE, directly from radio counters.

IV. RESULTS

Each of the five channel-count conditions (4, 8, 16, 32 and 64) was repeated three times, each run lasting 60 s with per-second sampling. Table II summarizes the outcome; throughout, we report medians for latency and RAN counters—whose means are contaminated by sporadic spikes and cumulative-counter artifacts—and means for the bitrate, loss, and overhead figures, whose run-to-run variation is negligible. The following subsections read the table along its three findings.

A. Scalability

The aggregate audio bitrate follows the ideal $N \times 64$ kbps reference almost exactly (Fig. 2 and Table II, column 2), with negligible run-to-run variation, and the RAN-side throughput counters exhibit the same linear scaling on both radio legs: the downlink to the Host UE and the uplink from the Artist UE scale in lockstep with the offered load up to roughly 6 Mbps at 64 channels. Because the application-layer demand and the independently measured radio-layer throughput scale in lockstep across the whole sweep, the offered load is accounted for at both layers. The RAN counters sit consistently above the IP-layer payload by a factor that decreases from $\approx 1.9 \times$ (4 ch)




Fig. 1. End-to-end system and measurement instrumentation. The concert audio (solid and dotted arrows) is generated at the artist workstation, carried uplink over the private 5G SA segment and delivered downlink to the receiving host. Two probes observe the flow at the application (1) and radio (2) layers. Dash-dotted arrows represent the internal flow of O-RAN.




Fig. 2. Aggregate audio bitrate (application layer) and per-UE downlink and uplink throughput (RAN layer) versus number of audio channels.

to $\approx 1.24 \times (64 \text{ ch})$, consistent with a fixed per-connection overhead (RTCP/ICE keepalives) plus per-packet lower-layer framing that dominate at low bitrate. Part of this residual gap may also be attributable to background traffic on the UE hosts unrelated to the audio streams. It is worth noting that the uplink is the constrained direction in this deployment: the TDD pattern of Table I allocates roughly 23% of the symbols to uplink transmission against 74% for downlink. The uplink counters nonetheless track the offered load as closely as the downlink ones across the whole sweep, indicating that the audio workload remains well within the available capacity even on the narrower leg.




Fig. 3. Median RTT (left) and jitter (right) versus number of audio channels.

B. Stable Latency

Latency does not grow with the offered load. As Fig. 3 shows, the median round-trip time stays at $\approx 34 \text{ ms}$ across the entire range (rising only marginally, to 36 ms , at 32 channels) and the median jitter stays between 5.0 and 8.2 ms , while the RAN contributes a sub-millisecond downlink RLC delay throughout ($137\text{--}449 \mu\text{s}$); the only load-dependent trend is a mild jitter increase at 64 channels, still well within the receiver playout budget. The two layers therefore tell the same story from opposite ends: a stable end-to-end RTT at the application, and a radio scheduler that adds negligible delay as seen from inside the RAN.

Packet loss remains negligible: although sporadic bursts occur, the per-second average stays below 0.2% even at 64 channels (Table II), and the bursts are uncorrelated with the channel count and non-monotonic in it (Fig. 4)—indeed the average loss is lower at 16 channels than at 8. The all-




Fig. 4. Per-second packet-loss distribution per condition.

tracks-at-once signature of these events is consistent with the synchronized burst structure described in Section III-B. The loss figures reported here are an upper bound on what a multitrack performance would experience.

C. Quantifiable cost of separate-track approach

Delivering each channel as an independent RTP stream is what makes per-channel spatialization and mixing possible at the receiver, but it is not free. Each 20 ms Opus frame (≈ 160 bytes) travels under a 40-byte RTP/UDP/IP header, so at 50 packets per second per track the protocol overhead adds ≈ 16 kbps per track—about 25% above the audio payload—and grows linearly to nearly 1 Mbps of headers alone at 64 channels. This is the structural price of the separate-track design relative to a single muxed stream, and it scales with the channel count rather than with the audio content.

D. Generality and extrapolation

One property bounds how far these results carry. The linear scaling makes channel count and per-stem bitrate interchangeable determinants of the aggregate load: any configuration within the tested ≈ 5 Mbps envelope—for instance 32 stems at 128 kbps—inherits the same stable RTT, jitter and RLC delay, whereas loads beyond it would require separate validation. The scope of this characterization is the transport of multichannel concert audio, from the encoder at the source to the receiving endpoint. Spatial rendering — Ambisonics encoding and decoding, per-source panning, and the associated computational cost — sits downstream of the delivery path and depends on the renderer and on the endpoint hardware rather than on the network. For the same reason, we do not report inter-channel playout alignment, since all N tracks share a single WebRTC PeerConnection and therefore traverse a common transport path.

V. CONCLUSIONS

We have shown that a private 5G Standalone network transparently sustains immersive multichannel concert audio delivered as real application-layer content rather than synthetic

traffic. Across a sweep from 4 to 64 simultaneous Opus tracks, the offered load scales linearly up to roughly 6 Mbps, while the median round-trip time and jitter remain load-independent and the radio access network contributes only sub-millisecond radio-link-control delay throughout. The result is established through a two-layer methodology in which application-level statistics and RAN-internal E2SM-KPM counters corroborate one another, and we have quantified the structural per-stream protocol overhead that the separate-track approach pays for its per-channel manipulability.

The broader significance lies in how little of the network this audio consumes. Even at 64 channels the workload remains around 6 Mbps with stable latency, jitter and loss—more than an order of magnitude below the point-cloud video workloads characterized over the same private 5G testbed in our companion studies, which demand hundreds of Mbps per user. The audio component is therefore effectively transparent to the network: adding high-channel-count immersive sound to a visual XR stream costs little incremental bandwidth and does not perturb its latency budget. Taken together, the audio feasibility established here and the point-cloud/video characterization of the companion work cover the two dominant media of an immersive concert over the same infrastructure, and lay the groundwork for a complete, jointly optimized audio–video–XR delivery system.

These findings should be read within the scope of the study. We consider one-to-many downstream distribution rather than bidirectional Networked Music Performance; a single artist source and a single receiving host; pre-recorded multitrack content built from independently encoded copies of a representative stem at 64 kbps per channel; and host workstations rather than a standalone headset endpoint. The measured transport latency is encouraging for interactive use, but the application-level playback buffering employed here is not yet compatible with the tighter closed-loop budgets of true Networked Music Performance and would need to be reduced accordingly.

Building on this basis, future work will pursue the integrated audio–video–XR experience end to end: combining the multichannel audio delivery characterized here with the immersive video component of the HEAT concert over a single 5G segment, and extending it to standalone headsets through the wearable 5G-to-USB gateway. Alongside this, we will address multi-performer, low-latency Networked Music Performance, closed-loop RAN control through purpose-built xApps to actively protect the audio flows, scaling to multi-UE and multi-listener configurations, and longer sessions, over which clock drift between endpoints may accumulate and affect inter-channel alignments.

REFERENCES

- [1] L. Turchet and M. Barthet, “An ubiquitous smart guitar system for collaborative musical practice,” *Journal of New Music Research*, vol. 48, pp. 352–365, Aug. 2019. [_eprint: https://doi.org/10.1080/09298215.2019.1637439](https://doi.org/10.1080/09298215.2019.1637439).
- [2] L. Turchet, C. Rinaldi, C. Centofanti, L. Vignati, and C. Rottondi, “5G-Enabled Internet of Musical Things Architectures for Remote Immersive Musical Practices,” *IEEE Open Journal of the Communications Society*, vol. 5, pp. 4691–4709, 2024.

- [3] L. Turchet, M. Lagrange, C. Rottondi, G. Fazekas, N. Peters, J. Østergaard, F. Font, T. Bäckström, and C. Fischione, "The Internet of Sounds: Convergent Trends, Insights, and Future Directions," *IEEE Internet of Things Journal*, vol. 10, pp. 11264–11292, July 2023.
- [4] L. Turchet, C. Fischione, G. Essl, D. Keller, and M. Barthet, "Internet of Musical Things: Vision and Challenges," *IEEE Access*, vol. 6, pp. 61994–62017, 2018.
- [5] E. Pilia, S. Porcu, G. Fadda, A. Palladino, M. Farina, L. Atzori, and M. Murrioni, "A Wearable 5G-to-USB Gateway for Standalone VR Headsets in 5G Networks," in *2026 International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, 2026. Accepted for publication.
- [6] 3GPP, "3GPP TR 26.928 Release 16."
- [7] M. Gapeyenko, V. Petrov, S. Paris, A. Marcano, and K. I. Pedersen, "Standardization of Extended Reality (XR) over 5G and 5G-advanced 3GPP new radio," *IEEE Network*, vol. 37, no. 4, pp. 22–28, 2023.
- [8] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges," *IEEE Communications Surveys & Tutorials*, vol. 25, no. 2, pp. 1376–1411, 2023.
- [9] C. Rottondi, C. Chafe, C. Allocchio, and A. Sarti, "An Overview on Networked Music Performance Technologies," *IEEE Access*, vol. 4, pp. 8823–8843, 2016.
- [10] J.-P. Cáceres and C. Chafe, "JackTrip: Under the hood of an engine for network audio," *Journal of New Music Research*, vol. 39, no. 3, pp. 183–187, 2010. [tex.eprint: https://doi.org/10.1080/09298215.2010.481361](https://doi.org/10.1080/09298215.2010.481361).
- [11] L. Gabrielli and S. Squartini, "Wireless networked music performance," in *Wireless networked music performance*, Springer, Singapore, 2016.
- [12] M. Sacchetto, P. Gastaldi, C. Chafe, C. Rottondi, and A. Servetti, "Web-based networked music performances via WebRTC: a low-latency PCM audio solution," *Journal of the Audio Engineering Society*, vol. 70, no. 11, pp. 926–937, 2022.
- [13] K. Tsioutas, Y. Thomas, F. Bistas, I. Barous, G. Xylomenos, and G. C. Polyzos, "Network music performance beyond 4G," in *2025 21st int. Wireless communications and mobile computing conf. (IWCMC)*, 2025.
- [14] X. Jiang, H. Shokri-Ghadikolaei, G. Fodor, E. Modiano, Z. Pang, M. Zorzi, and C. Fischione, "Low-latency networking: Where latency lurks and how to tame it," *Proceedings of the IEEE*, vol. 107, no. 2, pp. 280–306, 2019.
- [15] P. Ostan, C. Centofanti, M. Pezzoli, A. Bernardini, C. Rinaldi, and F. Antonacci, "Dynamic real-time Ambisonics order adaptation for immersive networked music performances," in *2025 33rd european signal processing conference (EUSIPCO)*, 2025.
- [16] M. Narbutt, J. Skoglund, A. Allen, M. Chinen, D. Barry, and A. Hines, "AMBIQUAL: a full reference objective quality metric for compressed ambisonic spatial audio," *Applied Sciences*, vol. 10, no. 9, p. 3188, 2020.
- [17] P. Hande, P. Tinnakornrisuphap, J. Damjanovic, H. Xu, M. Mondet, H. Lee, and I. Sakhnini, "Extended reality over 5G—standards evolution," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 6, pp. 1757–1771, 2023.
- [18] I. F. Akyildiz and H. Guo, "Wireless communication research challenges for extended reality (XR)," *ITU Journal on Future and Evolving Technologies*, vol. 3, no. 2, pp. 273–287, 2022.
- [19] J. Son, Y. Sanchez, C. Hellge, and T. Schierl, "Split rendering with L4S over 5G for latency critical interactive XR applications," *IEEE Communications Magazine*, vol. 62, no. 8, pp. 46–52, 2024.
- [20] O. S. Peñaherrera-Pulla, S. B. Damsgaard, B. Yanakiev, P. Mogensen, S. Fortes, and R. Barco, "Cloud VR on 5G: a performance validation in industrial scenarios," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 3641–3657, 2024.
- [21] A. Dominguez, B. da Costa Paulo, I. Burguera, I. Tamayo, A. Elosegi, S. Cabrero Barros, and M. Zorrilla, "A web-ready and 5G-ready volumetric video streaming platform: a platform prototype and empirical study," *IEEE access : practical innovations, open solutions*, vol. 14, pp. 34655–34675, 2026.
- [22] F. Kaltenberger, A. P. Silva, A. Gosain, L. Wang, and T.-T. Nguyen, "OpenAirInterface: Democratizing innovation in the 5G era," *Computer Networks*, vol. 176, p. 107284, 2020.
- [23] R. Schmidt, M. Irazabal, and N. Nikaein, "FlexRIC: An SDK for next-generation SD-RANs," in *Proc. 17th int. Conf. on emerging networking experiments and technologies (CoNEXT '21)*, pp. 411–425, 2021.
- [24] M. Murrioni, G.-M. Muntean, M. Montagud, and V. Popescu, "Toward interactive and distributed hybrid extended reality experiences," 2024.
- [25] G. Fiorina, M. Wildbaum, C. C. González, E. F. Pupo, A. Floris, S. Porcu, M. Murrioni, and R. Sotelo, "Estimating Quality of Experience in Multicast Point Cloud Streaming over 5G Networks," in *2025 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, pp. 1–6, June 2025. ISSN: 2155-5052.
- [26] N. Amarie, G. Fadda, V. Popescu, J.-M. Ghenta, M. Murrioni, and A. A. Simiscuka, "Live Feedback for Immersive Music Performances - A Case Study," in *ACM International Conference on Interactive Media Experiences Workshops (IMXw)*, pp. 26–29, SBC, June 2025.