Performance Analysis of Slicing on a 10-node 5G Architecture for Networked Music Performances

Luca Turchet
Dept. of Information Engineering and Computer Science
University of Trento
Trento, Italy
luca.turchet@unitn.it

Paolo Casari
Dept. of Information Engineering and Computer Science
University of Trento
Trento, Italy
paolo.casari@unitn.it

***Abstract*—Networked Music Performances (NMPs), where geographically displaced musicians play together over a data network, represent a challenging application for today's wireless communications. This is due to the stringent constraints on latency, throughput, and reliability that need to be obeyed in order to achieve a satisfactory quality of experience for the musicians. Slicing is a promising feature of 5G networks in the context of NMP applications, as it makes it possible to isolate the networking and computing resources allocated to NMP devices. However, the use of slicing has been scarcely investigated for the NMP context so far. Moreover, previous works focusing on NMPs over 5G involved up to 4 nodes. To bridge these gaps, we study 5G performance in support of NMPs involving an architecture with 10 nodes, both with and without slicing. Specifically, we focused on the assessment of the sole wireless link, as the measurements can be easily transferred to a realistic NMP architecture involving a wide area network. Our results show that, in the slicing condition, latency slightly increased due to the realistically different computing specifications of the MEC servers compared to those of the Core Network servers, whereas reliability slightly improved as expected.**

***Index Terms*—5G networks, slicing, networked music performances, Internet of Musical Things, latency, reliability.**

I. Introduction

The Internet of Musical Things (IoMusT) is an emerging field in communications and computing which refers to the extension of the Internet of Things (IoT) paradigm to the musical domain [1]. One of the central applications within the IoMusT is that of Networked Music Performances (NMPs), where geographically displaced musicians play together over wireless, wired, or hybrid networks via dedicated NMP systems [2], [3]. NMPs represent an active area of research, with an increasing body of literature focusing on the different disciplines of music perception, music technology and telecommunications and a marked focus on achieving practical results [4]–[9].

Relevant examples of NMP systems, either at the commercial or at the experimental level, are Aretousa [10], UNISON [11], fast-music [12], Elk LIVE [13], JackTrip [14], and LOLA [15]. While most NMP systems consisted of software applications executable on PCs and laptops in the past decades, advanced NMP technologies are now based on dedicated hardware platforms. These platforms leverage embedded systems that are optimized to reduce the delay of analog-to-digital and digital-to-analog audio signal conversions, as well as the delays due to audio processing and buffering [16].

Low-latency and highly-reliable communications are crucial to ensure a satisfying Quality of Experience (QoE) for musicians involved in an NMP session [17], [18]. Several perceptual studies have consistently shown that the maximum tolerable latency by displaced musician is about 30 ms (for a review the reader is referred to [3]). Moreover, audio transfers must be very reliable to prevent musicians from perceiving quality drops in the signal heard. In order to minimize latency, NMP systems typically rely on loss-tolerant protocols such as the User Datagram Protocol (UDP), and avoid optimizing the amount of bandwidth used as well as resorting to more reliable protocols coupled with audio compression methods. Due to such strict Quality of Service (QoS) requirements, NMP applications represent a challenge for wireless networks, and in particular for the current fifth generation (5G) of cellular networks [19], [20], the latest generation of mobile networks standardized by the 3rd Generation Partnership Project (3GPP). In general, scarce research has been conducted on interactive audio applications over 5G networks to date, with only a few studies focusing on such a topic with deployments in laboratory conditions [21] or actual testbeds [22].

A 5G end-to-end network (private or public) typically comprises three elements [23]:

1) *Core Network (CN)*: the central part of a network that provides services to users through the access network, and enables the transmission of IP packets to external networks such as the Internet.
2) *Radio Access Network (RAN)*: the network infrastructure, that includes radio base stations and bridges the connection between mobile radio network devices and the CN.
3) *User Equipment (UE)*: any device directly used by an end user to communicate. This includes mobile smartphone appliances, communication systems embedded in low-power edge devices, as well as massive IoT

This work received support from the MUR PNRR PRIN 2022 grant, prot. no. 2022CZWWKP, funded by Next Generation EU.

communication devices.

In addition, the 5G standard promoted by 3GPP has provisions for the following features:

- *Slicing*: the arrangement of a network into a set of logically separated, self-contained, and independent subnetworks, each of them dedicated to a different purpose or to services with different QoS requirements;
- *Multi-access Edge Computing (MEC)*: an architecture that provides cloud computing capabilities at the edge of the network (i.e., close to the base station), making it possible to reduce latency, ensure highly efficient network operations and service delivery, and improve the customer experience; notably, by deploying services and by caching content at the network edge, the CN is relieved from managing part of the traffic, making it possible to distribute computing resources to local tasks more effectively;
- *User Plane Function (UPF)*: a Virtual Network Function responsible for user packet routing and forwarding, packet inspections, and QoS handling; the UPF may be located on the MEC or on the CN.

Slicing, in particular, is a promising feature of 5G networks in the context of NMP applications, as it makes it possible to isolate the resources allocated to NMP devices, as opposed to sharing all radio resources among all mobile UEs. However, thus far the use of the 5G slicing feature has been scarcely investigated for the context of NMP systems. Moreover, the number of UEs involved in previous works focusing on NMPs over 5G has been limited, typically from 2 (e.g., see [24], [25]) to 4 (e.g., see [26]).

To bridge these gaps, in this paper we report a study assessing the 5G performance in supporting NMPs over an architecture with 10 UEs, with and without slicing. Specifically, we focused on the assessment of the sole wireless link, as the measurements can be easily transferred to a realistic NMP architecture involving a wide area network (WAN) by compounding them with transport delays and losses over the WAN.

II. Experimental setup

The 5G architecture deployed in our study was hosted at the premises of ZTE Italia Innovation & Research Center (ZIRC), located in L'Aquila (Italy), and was based on the architecture described in our previous study [26]. It comprised the following components.

User Equipment. The experimental setup involved 10 UEs. These consisted of a Customer Premise Equipment (CPE, i.e., a ZTE MC801A1 5G/WiFi/Ethernet router) connected to an audio/network interface device providing a peer-to-peer NMP system. Specifically we used the Elk LIVE NMP system by Elk, which is based on the low-latency audio operating system optimized for embedded systems reported in [13]. Each UE was associated to a laptop, through which the NMP session could be activated.

To automate the experiment sessions we created a Pure Data program, running on a laptop, that simulated the musical audio streams that would be produced by real musicians. Specifically, we used 10 audio signals that corresponded to recordings of 10 musicians playing together but recorded separately (namely two singers, two electric bass, two drums, two keyboard and two electric guitar players). Such audio streams were sent from a laptop running the Pure Data program to the 10 NMP devices, through an RME Fireface UFX II soundcard. The Elk Live system makes it possible to run multiple NMP sessions involving up to 5 musicians. Thus, the 10 NMP devices were allocated to 2 concurrent NMP sessions. Each device mixed the local input audio stream with those received from the other four devices. The resulting mixed signal could then be heard from headphones plugged in each device.

Each NMP device transmitted a stereo signals, sampled at 48 kHz. For each of the two audio channels, the device produces a protocol data unit comprising 64 audio samples at 16 bits/sample. The UDP protocol is employed at the transport layer without including any forward error correction or automatic repeat request scheme to protect the stream. As two audio channels are involved, the total protocol data unit (PDU) size is 272 bytes, and the packet transmission rate is one packet every $64/(48 \cdot 10^3) \approx 1.33$ ms. We measured that the minimum data rate required to transport all audio data seamlessly was approximately 15.5 Mbit/s for each NMP device in each of the two sessions. This led to a total offered traffic of 77.5 Mbit/s per session and of 155 Mbit/s considering both sessions.

Each NMP device, introduced a deterministic delay which amounts to 14.32 ms. This includes 0.5 ms for analog-to-digital and 0.5 ms for digital-to-analog conversions; ≈ 1.33 ms for the audio buffer utilized by the audio host for input and output (i.e., 64 samples at a sampling rate of 48 kHz); ≈ 10.66 ms for the jitter buffer. These settings left a latency budget of up to 15.68 ms to network transit in order to avoid exceeding the total latency tolerable by musicians (i.e., 30 ms).

Radio Access Network. In our setup, we used a ZTE QCell R8149 base station. This was placed on the ceiling, at a distance of about 3 m from the 10 UEs, which were located on top of an office table (see Fig. 1). Using ZTE's proprietary data rate metering software, we measured an available bandwidth of 1000 Mbit/s in downlink and 270 Mbit/s in uplink, nominally well above the requirements of the audio application.

5G Core Network. The CN hardware was located in the same building as the base station, about 10 m apart, and connected via a fiber optic cable. The CN included 8 servers, of which 3 were devoted to computing and network function hosting, including the UPF. A standard proportional fair scheduler was used, without any non-standard priority settings.

MEC. A ZTE ZXRAN U9003 MEC server was connected next to the base station through a 2-meter optical cable. It was configured to act as a relay of the audio packets traffic between the peers (i.e., a TURN server). To replicate real-world deployments, the MEC encompassed 5 servers, of which only one dedicated to the UPF deployment.

Slicing. The 5G system was configured to support two experimental conditions: without and with slicing. In the latter case, two slices were set up: the first was dedicated to the NMP service and was configured to guarantee higher resource scheduling priority; the second was configured for non-NMP communications. In the presence of slicing, the MEC hosted both the TURN server and the UPF, making it possible to avoid routing all traffic through the CN. Moreover, the MEC reserved resources for the audio streams of the UEs dedicated to the NMP. In the absence of slicing, both the NMP and the non-NMP services contended for the same network resources, thus sharing them at the same priority level.

III. Experiment

The evaluation procedure consisted in operating the IoMusT ecosystem in both the slicing and no-slicing conditions, during which the 10 NMP devices continuously transmitted audio packets to one another. Specifically, three measurement sessions, each lasting 5.30 min, were performed for each condition. Thanks to the logging system located in each NMP device, we measured the four metrics of interest in our analysis (namely latency, packet error ratio, missed packets, as well as the maximum number of consecutively missed packets). We computed such metrics over time periods of ≈ 2.33 s. Each time period contains 1750 packets, each carrying 64 audio samples. For each recording, we discarded the first and last 15 s in order to remove extra delays or imperfect synchronization effects due to the initial handshaking or the final disconnection of the devices. Finally, this yielded the observation of $\approx 225\,000$ packets sent by each NMP device in each session, for a total of $\approx 2\,250\,000$ packets analyzed in each session.

A. Results

Table I shows the results of the mean, standard deviation, minimum and maximum values measured on each NMP device across the four metrics of interest, both with and without slicing. Fig. 2 shows the evolution of the investigated metrics as recorded at one of the 10 NMP devices in the slicing and no-slicing conditions.

An analysis of variance was performed on different linear mixed effect models, one for each of the four metrics. For each model, we set the metric and the slicing condition as fixed factors, and the NMP device as the random factor. Regarding latency, a significant main effect was found for factor condition ($F(1, 37297) = 2601.6$, $p < 0.001$), showing that activating slicing introduced a significantly higher latency. As for the packet error ratio and number of missed packet, a significant main effect was found for factor condition ($F(1, 37263) = 12.51$, $p < 0.001$), showing that slicing led to a lower packet error ratio and fewer packets missed than the non-slicing condition. Regarding the maximum number of consecutive missed packets, a statistically significant effect was found for factor condition ($F(1, 37280) = 30.76$, $p < 0.001$), showing that slicing helped avoid long packet error bursts.

We applied Pearson's correlation tests to search for possible correlations between latency and the other three measures, both with and without slicing. For the slicing condition, significant correlations of very low strength were identified: for latency-packet error ratio, $r = 0.15$; for latency-missed packets, $r = 0.15$; for latency-max number of consecutive missed packets, $r = 0.08$; all were significant at $p < 0.001$. For the no-slicing condition significant correlations of low strength were identified: for latency-packet error ratio, $r = 0.29$; for latency-missed packets, $r = 0.29$; for latency-max number of consecutive missed packets, $r = 0.3$; all were significant at $p < 0.001$.

TABLE I
Results with and without slicing. The values include the NMP system's jitter buffer of ≈ 10.66 ms.

With Slicing

	Mean	SD	Min	Max
Latency (ms)	24.24	0.39	22.97	26.62
Packet error ratio	0.006	0.007	0	0.12
Missed packets	10.65	12.38	0	222
Max number of consecutive missed packets	2.52	4.15	0	63

Without Slicing

	Mean	SD	Min	Max
Latency (ms)	23.95	0.56	22.43	39.21
Packet error ratio	0.0064	0.009	0	0.22
Missed packets	11.3	16.84	0	393
Max number of consecutive missed packets	3.2	13.74	0	391

IV. Discussion and conclusions

The results reported in Section III-A show that slicing concretely helped dedicate specific resources to NMP traffic by scheduling NMP communications with higher priority, and thus guaranteeing slightly higher reliability. Nevertheless, latency increased in the presence of slicing (i.e., the MEC architecture led to higher latency than the CN servers). This result can be explained by the realistically lower computational resources available in MEC servers compared to CN servers. However, while statistically significant, the differences found between the slicing and no-slicing conditions for both latency and reliability metrics do not mark a strong change in the performance of the NMP system as a whole. Notably, in our setup we focused on the sole wireless link, without involving a WAN. When WAN routing delays are taken into account, the slightly higher computation delay of MEC servers would still prevent much longer routing delays over geographically displaced entry points of an operator's core network. In this case, the benefits of a slicing-based 5G architecture would be even stronger.

Our findings show that our setup, which involved a single base station, could effectively support two NMP sessions with a total of 10 nodes, meaning that 10 musicians could

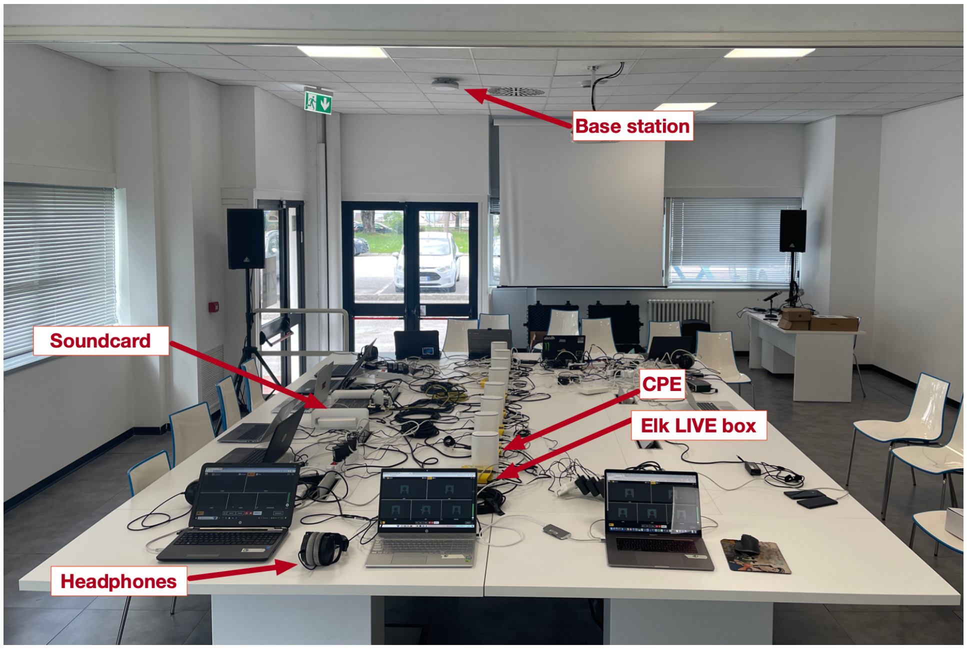

Fig. 1. A picture of the experimental setting of the 5G architecture, showing the base station (on the ceiling), the 10 CPEs, the 10 Elk LIVE NMP devices, headphones, the soundcard, and the laptops.

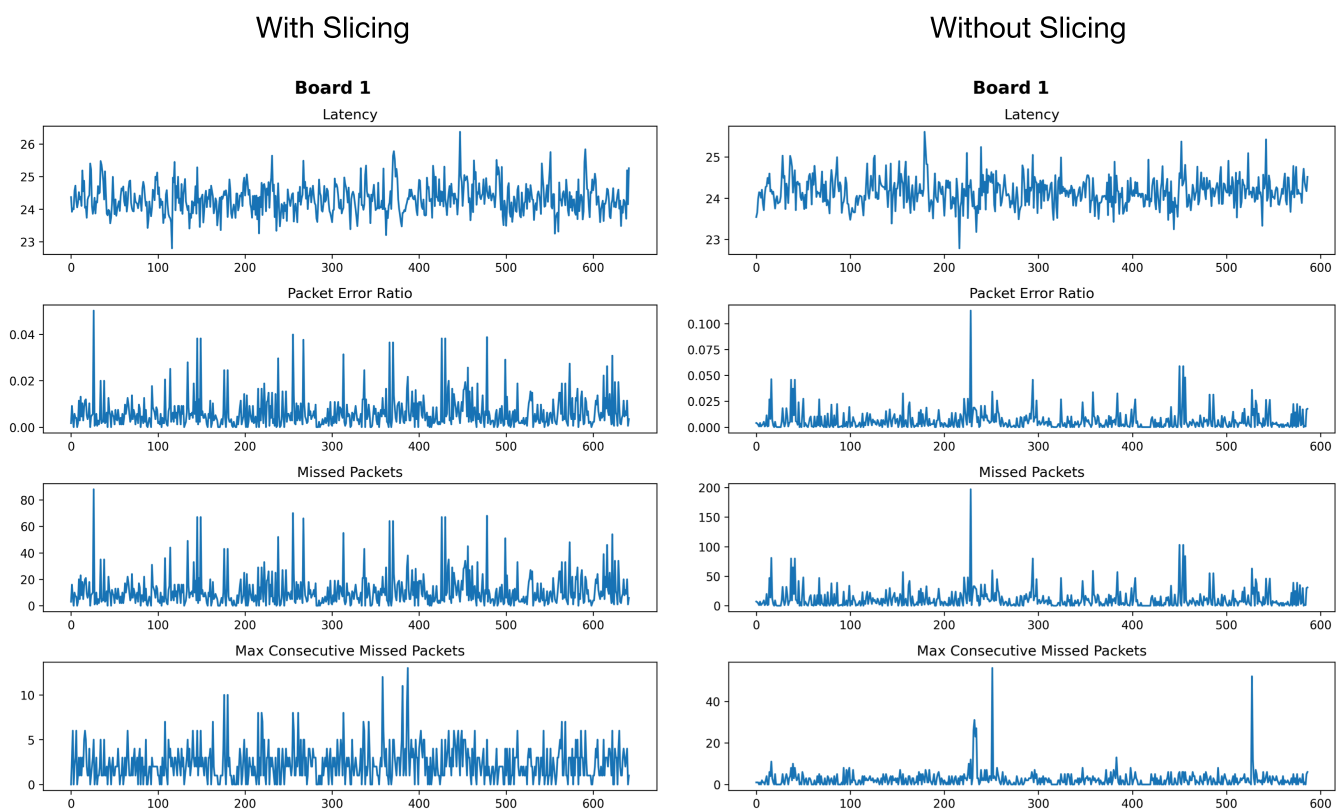

Fig. 2. Evolution of the four performance metrics recorded at one of the 10 boxes, over a period of 5 minutes, for the condition with slicing (left) and without it (right). Notice the different y-axis scale on the left and right panels for what concerns the three reliability metrics.

take advantage of the NMP system to perform together. The average latency of about 24 ms left a time budget for the WAN contribution of about 6 ms, in order to avoid exceeding the perceptually-defined threshold of 30 ms. The fact that only 6 ms are left for the WAN contribution suggests that improvements on 5G networks are needed to properly support NMP sessions and in general IoMusT application, as recently highlighted by the authors of the study in [27]. Notably, in our setup we used a jitter buffer of 10.66 ms. This could be reduced to increase the time budget for the WAN, but at the cost of a lower reliability and resilience to delay variations. The performance decrease across reliability metrics could be compensated by effective packet loss concealment methods capable of reconstructing the missing signal with zero latency [28]–[32]. This is a matter still under investigation due to the complexities inherent in the issue of maximizing perceptual fidelity while minimizing processing latency and computational demands.

Latency and reliability results were found to be uncorrelated irrespective of slicing conditions. This finding suggests that these two key performance indicators are driven by different root causes. Notably, such a finding parallels those of other previous studies assessing the performance of 5G systems in supporting NMP performances [24]–[26], [33].

It is worth noticing that this study presents some limitations. In the first place, we conducted only three session per condition, which lasted just 5 minutes and were performed on the same day. Conducting more extensive and longitudinal tests would make it possible to generalize our findings to different scenarios. Second, due to the limited equipment availability, we could only test an architecture with 10 nodes, without involving a concurrent background traffic. This could provide further insight on slicing performance in support of IoMusT applications under more realistic worst case conditions. Third, musicians were simulated in our experiments. The involvement of human musicians would corroborate the objective findings reported here with subjective results. Finally, our study focused on a private 5G architecture, without involving a WAN, which would represent a more realistic scenario and is the object of future investigations.

In future work, we plan to conduct additional performance assessment experiments, involving an architecture with a higher number of NMP devices, as well as conditions with or without concurrent UDP traffic of different intensity. This will allow us to assess the limits of 5G systems in supporting IoMusT settings with a massive number of nodes under a same radio cell. Furthermore, we plan to conduct such tests with musicians.

Acknowledgment

We acknowledge the support of the MUR PNRR PRIN 2022 grant, prot. n. 2022CZWWKP, funded by Next Generation EU. The authors are grateful to Daniele Ciambrone and Domenico Puntillo (ZTE), and to Stefano Zambon (Elk), for their help and support.

References

[1] L. Turchet, C. Fischione, G. Essl, D. Keller, and M. Barthet, "Internet of Musical Things: Vision and Challenges," *IEEE Access*, vol. 6, pp. 61 994–62 017, 2018.

[2] L. Gabrielli and S. Squartini, *Wireless Networked Music Performance*. Springer, 2016.

[3] C. Rottondi, C. Chafe, C. Allocchio, and A. Sarti, "An overview on networked music performance technologies," *IEEE Access*, vol. 4, pp. 8823–8843, 2016.

[4] E. Lakiotakis, C. Liaskos, and X. Dimitropoulos, "Improving networked music performance systems using application-network collaboration," *Concurrency and Computation: Practice and Experience*, vol. 31, no. 24, p. e4730, 2019.

[5] P. Cairns, H. Daffern, and G. Kearney, "Immersive network music performance: Design and practical deployment of a system for immersive vocal performance," in *Audio Engineering Society Convention 149*. Audio Engineering Society, 2020.

[6] L. Comanducci, M. Buccoli, M. Zanoni, A. Sarti, S. Delle Monache, G. Cospito, E. Pietrocola, and F. Berbenni, "Investigating networked music performances in pedagogical scenarios for the intermusic project," in *IEEE Conference of Open Innovations Association (FRUCT)*, 2020, pp. 119–127.

[7] 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.

[8] L. Comanducci, "Intelligent networked music performance experiences," in *Special Topics in Information Technology*. Springer, Cham, 2023, pp. 119–130.

[9] K. Tsioutas and G. Xylomenos, "Assessing the Effects of Delay to NMP via Audio Analysis," *SN Computer Science*, vol. 4, no. 2, pp. 1–13, 2023.

[10] K. Tsioutas, G. Xylomenos, and I. Doumanis, "Aretousa: A competitive audio streaming software for network music performance," in *Audio Engineering Society Convention 146*. Audio Engineering Society, 2019.

[11] C. Werner and R. Kraneis, "Unison: a novel system for ultra-low latency audio streaming over the internet," in *2021 IEEE 18th Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, 2021, pp. 1–4.

[12] A. Carôt, C. Hoene, H. Busse, and C. Kuhr, "Results of the fast-music project—five contributions to the domain of distributed music," *IEEE Access*, vol. 8, pp. 47 925–47 951, 2020.

[13] L. Turchet and C. Fischione, "Elk Audio OS: an open source operating system for the Internet of Musical Things," *ACM Transactions on the Internet of Things*, vol. 2, no. 2, pp. 1–18, 2021.

[14] J. 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.

[15] C. Drioli, C. Allocchio, and N. Buso, "Networked performances and natural interaction via lola: Low latency high quality a/v streaming system," in *International Conference on Information Technologies for Performing Arts, Media Access, and Entertainment*. Springer, 2013, pp. 240–250.

[16] R. Moscatelli, K. Stahel, R. Kraneis, and C. Werner, "Why real-time matters: Performance evaluation of recent ultra-low latency audio communication systems," in *2024 IEEE 21st Consumer Communications & Networking Conference (CCNC)*. IEEE, 2024, pp. 77–83.

[17] R. Hupke, D. Jan, N. Werner, and J. Peissig, "Latency and quality-of-experience analysis of a networked music performance framework for realistic interaction," in *Audio Engineering Society Convention 152*. Audio Engineering Society, 2022, pp. 1–10.

[18] K. Tsioutas, G. Xylomenos, and I. Doumanis, "An empirical evaluation of qome for nmp," in *2021 11th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*. IEEE, 2021, pp. 1–5.

[19] F. Cheli and S. Giordano, "Service parameters identification for adaptive networked music performance," in *2022 Global Information Infrastructure and Networking Symposium*. IEEE, 2022, pp. 94–98.

[20] L. Vignati, G. Nardini, M. Centenaro, P. Casari, S. Lagén, B. Bojovic, S. Zambon, and L. Turchet, "Is Music in the Air? Evaluating 4G and 5G Support for the Internet of Musical Things," *IEEE Access*, 2024.

[21] C. Rinaldi, F. Franchi, A. Marotta, F. Graziosi, and C. Centofanti, "On the exploitation of 5G multi-access edge computing for spatial audio in cultural heritage applications," *IEEE Access*, vol. 9, pp. 155 197–155 206, 2021.

[22] F. Martusciello, C. Centofanti, C. Rinaldi, and A. Marotta, "Edge-Enabled Spatial Audio Service: Implementation and Performance Analysis on a MEC 5G Infrastructure," in *2023 4th International Symposium on the Internet of Sounds*. IEEE, 2023, pp. 1–8.

[23] L. Bonati, M. Polese, S. D'Oro, S. Basagni, and T. Melodia, "Open, programmable, and virtualized 5g networks: State-of-the-art and the road ahead," *Computer Networks*, vol. 182, p. 107516, 2020.

[24] A. Carôt, M. Dohler, S. Saunders, F. Sardis, R. Cornock, and N. Uniyal, "The world's first interactive 5G music concert: Professional quality networked music over a commodity network infrastructure," in *Proc.17th Sound and Music Computing Conf.*, 2020, pp. 407–412.

[25] J. Dürre, N. Werner, S. Hämäläinen, O. Lindfors, J. Koistinen, M. Saarenmaa, and R. Hupke, "In-depth latency and reliability analysis of a networked music performance over public 5G infrastructure," in *Audio Engineering Society Convention 153*. Audio Engineering Society, 2022.

[26] L. Turchet and P. Casari, "Latency and reliability analysis of a 5G-enabled Internet of Musical Things system," *IEEE Internet of Things Journal*, vol. 11, no. 1, pp. 1228–1240, 2024.

[27] L. Mohjazi, B. Selim, M. Tatipamula, and M. A. Imran, "The journey towards 6g: A digital and societal revolution in the making," *arXiv preprint arXiv:2306.00832*, 2023.

[28] M. Fink and U. Zölzer, "Low-delay error concealment with low computational overhead for audio over ip applications." in *Proceedings of the International Conference on Digital Audio Effects*, 2014, pp. 309–316.

[29] P. Verma, A. I. Mezza, C. Chafe, and C. Rottondi, "A deep learning approach for low-latency packet loss concealment of audio signals in networked music performance applications," in *2020 27th Conference of Open Innovations Association (FRUCT)*. IEEE, 2020, pp. 268–275.

[30] A. I. Mezza, M. Amerena, A. Bernardini, and A. Sarti, "Hybrid packet loss concealment for real-time networked music applications," *IEEE Open Journal of Signal Processing*, 2023.

[31] L. Vignati, S. Dallona, and L. Turchet, "PLC Testbench: a modular tool for the study and comparison of audio Packet Loss Concealment algorithms," in *2023 4th International Symposium on the Internet of Sounds*. IEEE, 2023, pp. 1–8.

[32] M. Sacchetto, C. Rottondi, and A. Bianco, "Implementation and optimization of burg's method for real-time packet loss concealment in networked music performance applications," *Personal and Ubiquitous Computing*, pp. 1–17, 2024.

[33] L. Turchet and P. Casari, "Assessing a Private 5G SA and a Public 5G NSA Architecture for Networked Music Performances," in *Proceedings of the 4th International Symposium on the Internet of Sounds*. IEEE, 2023, pp. 1–6.