# O-RIS-ing: Evaluating RIS-Assisted NextG Open RAN

Maria Tsampazi\*, Michele Polese\*, Falko Dressler<sup>§</sup>, Tommaso Melodia\*

\*Institute for the Wireless Internet of Things, Northeastern University, Boston, MA, U.S.A.

E-mail: {tsampazi.m, m.polese, t.melodia}@northeastern.edu

<sup>§</sup>School of Electrical Engineering and Computer Science, TU Berlin, Germany

E-mail: {dressler}@ccs-labs.org

Abstract—Reconfigurable Intelligent Surfaces (RISs) pose as a transformative technology to revolutionize the cellular architecture of Next Generation (NextG) Radio Access Networks (RANs). Previous studies have demonstrated the capabilities of RISs in optimizing wireless propagation, achieving high spectral efficiency, and improving resource utilization. At the same time, the transition to softwarized, disaggregated, and virtualized architectures, such as those being standardized by the O-RAN ALLIANCE, enables the vision of a reconfigurable Open RAN. In this work, we aim to integrate these technologies by studying how different resource allocation policies enhance the performance of RISassisted Open RANs. We perform a comparative analysis among various network configurations and show how proper network optimization can enhance the performance across the Enhanced Mobile Broadband (eMBB) and Ultra Reliable and Low Latency Communications (URLLC) network slices, achieving up to  $\sim 34\%$ throughput improvement. Furthermore, leveraging the capabilities of OpenRAN Gym, we deploy an xApp on Colosseum, the world's largest wireless system emulator with hardware-in-theloop, to control the Base Station (BS)'s scheduling policy. Experimental results demonstrate that RIS-assisted topologies achieve high resource efficiency and low latency, regardless of the BS's scheduling policy.

Index Terms—O-RAN, Open RAN, Wireless Network Emulator, Reconfigurable Intelligent Surfaces, Resource Allocation

#### I. INTRODUCTION

The Open RAN paradigm for O-RAN networks is pivotal in transforming cellular network design. Grounded in the principles of disaggregation, virtualization, and programmable radios, components are interconnected through open, standardized interfaces, ensuring multi-vendor interoperability. This approach enables the deployment of flexible architectures built on cloud-native principles. The reconfigurability of the RAN is further enhanced through RICs, streaming RAN telemetry and implementing control policies via a closed-loop mechanism [1]. Both near- and non-real time RICs, deployed at the RAN's edge, leverage xApps and rApps to optimize network management and resource allocation. This functionality further advances the integration of Artificial Intelligence (AI) for the RAN and AI on the RAN, fostering the development of intelligent and adaptive RAN architectures [2].

At the same time, RISs are envisioned to be an integral part of NextG cellular architectures, enabling smart propagation environments that can be controlled through programmable Radio Frequency (RF) components [3]. Thanks to their reconfigurability, RISs can manipulate electromagnetic waves, reshaping wireless propagation to optimize spectrum utilization and efficiency, particularly in multi-operator networks [4]. This reconfiguration capability aligns with O-RAN's objectives of developing adaptable and resource-efficient RANs.

Intelligent control loops with xApps in O-RAN have attracted widespread attention from the research community [5]. In [6], the authors explore the impact of xApps embedding Deep Reinforcement Learning (DRL) agents to manage the Base Station (BS)'s slicing and scheduling control policies across various network deployments on the Colosseum testbed. In addition, several efforts have been conducted in the optimization of wireless propagation with RISs. Some previous works [7]–[9] have focused on the use of RISs and constructive beams to achieve coherent signal construction at the BS, for maximizing the achieved Received Signal Strength Indicator (RSSI). Finally, a system-level experimental evaluation of RISenabled deployments and their cross-layer optimization, leveraging the full protocol stack on the Colosseum testbed has been discussed in [10].

# A. Contributions and Outline

However, limited research has focused on bridging O-RAN and RISs at a system level. In this study, we aim to address this research gap by evaluating network performance across various RIS-assisted topologies using the Colosseum Open RAN testbed [11]. We consider an Open RAN delivering services to the eMBB and URLLC network slices. Using OpenRAN Gym [5]—an open-source framework for experimentation in O-RAN-we deploy this Open RAN on the Colosseum network emulator and manage it through an xApp that controls the BS's scheduling policy. By investigating various spectrum allocation policies and RIS deployments, we demonstrate how network performance and resource efficiency are impacted differently for each slice. With this work, we aim to complement prior studies [10] by providing a comprehensive evaluation of RISassisted Open RANs on an O-RAN-compliant testbed. We believe that our findings and insights will contribute to a deeper understanding of the integration of RISs in NextG Open RANs.

The remainder of this paper is organized as follows. Section II describes the system model. Section III presents the optimization framework for shaping wireless propagation with RISs, while Section IV discusses the resource allocation strategies. Section V details the experimental setup, while Section VI discusses the experimental results. Finally, Section VII draws our conclusions and outlines directions for future work.

This article is based upon work partially supported by the U.S. National Science Foundation under grants CNS-1925601 and CNS-2112471.

# II. SYSTEM MODEL

#### A. Open RAN Framework

We investigate an Open RAN system where User Equipments (UEs) generate traffic with diverse profiles, which is classified into the eMBB and URLLC network slices. We leverage an xApp tasked with reconfiguring the BS's Medium Access Control (MAC)-layer scheduling policies. The xApp selects a dedicated scheduling profile among multiple resource allocation algorithms (i.e., Round Robin (RR), Waterfilling (WF), and Proportionally Fair (PF)), influencing how Physical Resource Blocks (PRBs) are internally allocated to the UEs of each slice [6]. It is noted that for the RF environment we consider various RIS-assisted network deployments. In Fig. 1, we outline the reference architecture considered in our work.

# B. Channel Model

We consider a topology consisting of a RIS, a BS, and a set of UEs, denoted by  $I = \{1, \ldots, i, \ldots, |I|\}$ , which establish direct communication with the BS. The number of elements on the RIS is given as |M| and their set is defined as  $M = \{1, \ldots, m, \ldots, |M|\}$ . The first RIS element, i.e., m = 1, with its respective coordinates, i.e.,  $(x_R, y_R, z_R)$  [m], is used as a reference point in the following arithmetic calculations, while for the BS, its coordinates are given as follows  $(x_A, y_A, z_A)$  [m]. The channel gain corresponding to the direct link between a UE i and the BS is given as  $h_{iA}$ , the channel gain between a UE i and the RIS is defined as  $h_{iR}$ , where  $\mathbf{\hat{h}}_{iR} = [|h_{iR,1}|e^{j\omega_1}, \dots, |h_{iR,m}|e^{j\omega_m}, \dots, |h_{iR,|M|}|e^{j\omega_{|M|}}]^T,$ and finally, the channel gain of the reflected link from the RIS to the BS is denoted as  $h_{RA}$ . The diagonal phase-shift matrix is defined as  $\Theta = \text{diag}(e^{j\theta_1}, \dots, e^{j\theta_m}, \dots, e^{j\theta_{|M|}})$ , while each RIS element's phase shift is  $\theta_m \in [0, 2\pi], \forall m \in M$ , and the cascaded channel is therefore given as  $\mathbf{h}_{RA}^{H} \Theta \mathbf{h}_{iR}$ . Finally, the steering vector for the UE *i* to the RIS link is given as  $\mathbf{h}_{iR}^{\prime\prime\prime} = [1, e^{-j\frac{2\pi}{\lambda}d\phi_{iR}}, \dots, e^{-j\frac{2\pi}{\lambda}(|M|-1)d\phi_{iR}}]^T$ , while the steering vector for the RIS to the BS link is denoted as  $\mathbf{h}_{RA}^{\prime\prime\prime} = [1, e^{-j\frac{2\pi}{\lambda}d\phi_{RA}}, \dots, e^{-j\frac{2\pi}{\lambda}(|M|-1)d\phi_{RA}}]^T$ . In the above calculations, d is the Euclidean distance,  $\lambda[m]$  is the carrier wavelength and d[m] is the antenna separation. Additionally,  $\phi_{RA}$  is the cosine of the signal's Angle of Departure (AoD) for the RIS to the BS link, and  $\phi_{iR}$  is the Cosine of the signal's Angle of Arrival (AoA) for the UE *i* to RIS link. It is noted that we consider the links between the UEs and the BS to be Non-Line-of-Sight (NLOS), while the channels connecting the UEs to the RIS and the RIS to the BS are assumed to be Lineof-Sight (LOS). These transmission conditions are chosen to ensure that the cascaded channel formed by the RIS provides a sufficiently strong link for redirecting the UEs' signals to the BS. Such placement can represent a maneuverable flying RIS mounted on an Unmanned Aerial Vehicle (UAV), capable of dynamically establishing LOS to UEs that would otherwise encounter NLOS relative to a fixed-position BS.

Lastly, both the BS and the UEs are equipped with singleantenna omnidirectional Transmitters (TXs)/Receivers (RXs), with their center frequency set to 5.9 GHz, where RISs are considered for deployment [12]. For the generation of the wireless



Fig. 1. Reference O-RAN testing architecture for RIS-enabled deployments, focusing on the Sched xApp Case Study, as described in Section VI-C.

links, we leverage the capabilities of the QUAsi Deterministic RadIo channel GenerAtor (QuaDRiGa)'s [13] statistical raytracing approach, which is based on channel sounding and geometry-based channel modeling. We observe the channels every 0.1 s and select the 3GPP\_38.901\_UMa scenario for the channel generator. A number of Multipath Components (MPCs) (denoted as |N|) are produced for each generated link, with the channel generation procedure based on the following flow adopted from [10]:

- 1) Generate scenario-specific MPCs (denoted as  $h'_{xy}$ ) for each link, leveraging QuaDRiGa.
- 2) Combine the MPCs coherently to compute the channel gain for every link, following Eq. (1):

$$\mathbf{h}_{xy}'' = \sum_{j=1}^{|N|} \left| \mathbf{h}_{xy}' \right| \cdot e^{j\varphi_{xy}'},\tag{1}$$

The MPC gain is denoted as  $h'_{xy}$ , where  $x \in \{i, R\}, y \in \{R, A\}$ , and  $|h'_{xy}|, \varphi'_{xy}$  are the respective MPC magnitude and phase.

- Calculate the average of the channel gains (i.e., h<sup>"</sup><sub>xy</sub>) over 100 different channel realizations to determine the link's final channel gain (i.e., h<sub>iA</sub>).
- Multiply the computed channels by their corresponding steering vectors, denoted as h<sup>'''</sup><sub>RA</sub> and h<sup>'''</sup><sub>iR</sub>, to obtain the respective final channel gains (i.e., h<sub>RA</sub>, h<sub>iR</sub>).
- 5) Calculate the overall channel power gain for the UE *i* to BS link as follows  $G_i = |h_{iA} + \mathbf{h}_{BA}^H \Theta \mathbf{h}_{iR}|^2$ .

#### III. OPTIMIZING WIRELESS PROPAGATION WITH RIS

In this work, the wireless environment is not treated as an uncontrollable entity but is instead subject to optimization. For this purpose, we leverage the capabilities of RISs to reshape the wireless propagation by effectively adapting the phase shifts of the RIS elements. For these reasons, we formulate the following optimization problem defined in Eqs. (2a)-(2b), as follows:

$$\max_{\boldsymbol{\theta}} \sum_{i=1}^{|I|} |h_{iA} + \mathbf{h}_{RA}^{H} \boldsymbol{\Theta} \mathbf{h}_{iR}|^2$$
(2a)

s.t. 
$$0 \le \theta_m \le 2\pi, \forall m \in M.$$
 (2b)

It is noted that the solution to the aforementioned optimization problem, which takes place on the BS's side, is the effective

ş

phase shift vector of the RIS elements, denoted as  $\theta^*$ . In order to maximize the overall channel power gain of the UEs, the BS only needs to calculate the effective phase shifts of the RIS elements (i.e.,  $\theta^*$ ). Given the non-convexity of the aforementioned optimization problem, computing a globally optimal solution is challenging. To derive a closed-form phase-shift solution, we aim to achieve coherent signal alignment for the incoming signals from different transmission paths at the BS. Indeed, the UE's channel power gain is maximized when the direct signal and the signal reflected by the RIS are perfectly aligned and coherently combined at the BS's RX [7], [8]. This condition is generally satisfied when the phase shifts of the direct and cascaded signals are aligned as  $\angle h_{iA} = \angle (\mathbf{h}_{RA}^H \Theta \mathbf{h}_{iR})$ . Therefore, in the single-UE case there's an  $1 \times |M|$  phase-shift vector  $\theta^* = \angle \mathbf{v}$ , where the optimal phase shifts of the RIS elements are given as follows:

$$\theta_m^* = \angle h_{iA} + \omega_m + \frac{2\pi}{\lambda} d(m-1)\phi_{RA}, \forall m \in M.$$
 (3)

In the multi-UE case, the solution to the maximization problem is also a linear combination of the phase shifts of the RIS elements, with the reflection-coefficient vectors  $\mathbf{v}_i$  being distinct for each UE. Therefore, there exists a distinct reflectioncoefficient vector  $\mathbf{v}_i = [v_{i,1}, \ldots, v_{i,|M|}] \in \mathbb{C}^{|M| \times 1}$  for each UE *i* that maximizes its channel power gain. To determine the set of suitable weights for each UE that will maximize the aggregate channel power gain across all UEs as defined in Eq. (2a), we adopt the research methodologies proposed in [9], [14]. For each UE *i*, we define an appropriate weight factor  $w_i \in [0, 1]$  and compute the linear combination of the overall UEs' reflection-coefficients  $\mathbf{v}_i$  as follows:

$$\mathbf{v} = \sum_{i=1}^{|I|} w_i \mathbf{v}_i,\tag{4}$$

such as the weight factors satisfy the condition  $\sum_{i=1}^{|I|} w_i = 1$ . In the multi-UE case, the goal is to determine the optimal values of the weight factors (i.e.,  $w_i$ ) for the reflection-coefficient vector (i.e., **v**) of Eq. 4, which maximize the overall channel power gain of the UEs, as defined in Eq. (2a). Therefore, in order to compute the effective RIS elements phase shifts, the optimization problem defined in Eqs. (2a)-(2b) is formulated as follows in Eqs. (5a)-(5c):

$$\max_{\mathbf{w}} \sum_{i=1}^{|I|} |h_{iA} + \mathbf{h}_{RA}^{H} \boldsymbol{\Theta} \mathbf{h}_{iR}|^2$$
(5a)

s.t. 
$$0 \le w_i \le 1, \forall i \in I$$
 (5b)

$$\sum_{i=1}^{|I|} w_i = 1 \tag{5c}$$

Through backward induction [9], the RIS reflection matrix  $\Theta = \text{diag}(e^{j \angle \mathbf{v}})$  is computed to optimize the phase alignment. This approach ensures that the final RIS configuration  $\Theta$  maximizes both the channel gain of each UE and the total channel power gain across all UEs. Given that  $\mathbf{w} = [w_1, \ldots, w_i, \ldots, w_{|I|}]$  is the vector of the UEs' assigned weight

factors, the optimization problem described in Eqs. (5a)-(5c) consists of a non-negative linear objective function and a set of constraints, which can be intuitively optimized to derive the optimal weights  $\mathbf{w}^*$ . The derivation of the optimal weights  $\mathbf{w}^*$  leads to an effective RIS elements' reflection-coefficient vector  $\mathbf{v}^* = [v_1, \ldots, v_m, \ldots, v_{|M|}]$ , which subsequently determines the corresponding phase shifts of the RIS elements, expressed as  $\boldsymbol{\theta}^* = [\theta_1, \ldots, \theta_m, \ldots, \theta_{|M|}]$ .

# **IV. RESOURCE ALLOCATION STRATEGIES**

We consider three different case studies, encompassing both single and multi-UE scenarios, which are evaluated under various resource management policies in terms of PRB allocation and scheduling profile selection.

**Case Study A.** Following the requirements outlined in [12], which specify that existing RIS panels can host a minimum of 10 elements, we evaluate three network configurations, where the number of RIS elements increases by orders of magnitude. We consider a single UE allocated to the eMBB network slice, which is assigned  $\sim 1/3$  of the available bandwidth resources. The Channel State Information (CSI) is varied to evaluate the contribution of RIS to network performance given the same resource availability, which is considered fixed. The results of this analysis are provided in Section VI-A.

**Case Study B.** Considering the trade-off between hosting thousands of RIS elements and maintaining operational simplicity, we focus on a RIS with 100 elements, serving multiple UEs, which are distributed across the eMBB and URLLC network slices. By dividing the available spectrum between the two, we aim to demonstrate how a RIS can significantly enhance network performance, improve UE satisfaction, and increase PRB utilization based on the capacity demands of each slice. The findings of this investigation are provided in Section VI-B.

**Case Study C.** This case study focuses on an O-RAN infrastructure where an xApp controls the scheduling profile selection of the BS's Distributed Unit (DU). The xApp updates the scheduling profile of the BS's scheduler through a control loop by alternating its policy given the following options: RR, WF, and PF. A multi-UE use case is considered, where eMBB UEs are provisioned with 90% of the available bandwidth, while URLLC UEs are allocated the remaining 10% of the spectrum resources. With this case study, we aim to study the impact of an xApp frequently changing the scheduling profile from fair scheduling (i.e., RR) to spectral efficiency-driven (i.e., WF), or balanced scheduling (i.e., PF), analyzing its effect on resource allocation and overall network performance in both RIS-assisted and non-RIS-assisted topologies. The results of this exploration are provided in Section VI-C.

#### V. EXPERIMENTAL SETUP

To experimentally evaluate the RIS-assisted channels within an O-RAN ecosystem, we leverage the capabilities of Open-RAN Gym, an experimental toolbox for the end-to-end development, implementation, and testing of diverse solutions including, but not limited to, AI/Machine Learning (ML) applications—within O-RAN. The software offers capabilities such as RAN and core network deployments using the srsRAN [15] protocol stack. It supports large-scale data collection, testing, and fine-tuning of RAN functionalities by integrating open Application Programming Interfaces (APIs) for slicing and scheduling control, as well as Key Performance Measurements (KPMs) collection. Additionally, the software incorporates an O-RAN-compliant control architecture for executing xApps in the near-real-time RIC. Through the E2 interface, which bridges the RAN and RIC, along with its Service Models (SMs) [1], the streaming of KPMs from the RAN and the execution of control actions by the xApps are enabled.

We deploy OpenRAN Gym on Colosseum, which features 128 Standard Radio Nodes (SRNs) consisting of pairs of Dell PowerEdge R730 servers and NI Universal Software Radio Peripheral (USRP) X310 Software-defined Radios (SDRs), supporting large-scale experimentation in diverse network deployments. The testbed's Massive Channel Emulator (MCHEM) component emulates the wireless environment by leveraging Field Programmable Gate Array (FPGA)-based Finite Impulse Response (FIR) filters. These filters simulate RF conditions, including path loss, fading and attenuation based on models created through various frameworks (e.g., ray-tracing software, analytical models, or real-world measurements). Similarly, the Colosseum Multi-Generator (MGEN) TCP/UDP traffic generator [16] emulates a variety of network traffic profiles (e.g., multimedia content) and demand distributions (e.g., Poisson, periodic).

We consider a cellular network comprising one BS and up to 5 UEs distributed across 2 network slices. These are: (i) eMBB, for high-traffic scenarios involving multimedia content and streaming applications; and (ii) URLLC, designed for timecritical applications such as vehicle coordination in Cellular Vehicle-to-Everything (C-V2X) environments. The maximum bandwidth considered in this work is is set to 10 MHz (i.e., 50 PRBs) and is divided among the 2 slices. Traffic is slice-based and generated according to the following specifications: eMBB UEs request a constant bitrate of 4 Mbps, while URLLC UEs generate Poisson traffic at 89.3 kbps. In terms of physical deployment, the BS is placed at the point ( $x_A = 25$ ,  $y_A = 50$ ,  $z_A = 25$ ) [m] of the three-dimensional space. Their dis-

The second secon

Fig. 2. Path Gains under varying numbers of RIS elements.

tance from the RIS is given as  $d_{i,R} = [20, 27, 37, 58, 66]$  [m] and all are uniformly distributed around the RIS's reference point, denoted as  $(x_R = 30, y_R = 40, z_R = 20)$  [m].

We install the generated links in Colosseum. The Colosseum scenario includes the path gain (i.e., a complex FIR coefficient) and the Time-of-Arrival (ToA) values, which are are calculated based on the Euclidean distance between the active entities (i.e., BS, RIS, and UEs), using the speed of light. These values are used to generate the time-variant Channel Impulse Response (CIR) for each node pair in the network.

#### VI. EXPERIMENTAL EVALUATION<sup>1</sup>

In Fig. 2, we present the path gains for each UE, i, with and without a RIS, generated through the channel modeling process and the optimization procedure as described in Sections II-B and Section III, respectively. All the UEs are sorted in descending order, such that the UE with the lowest index corresponds to the UE with the best channel, while the UE with the highest index corresponds to the UE with the worst. Based on the reported results, the performance with as few as 10 RIS elements is nearly identical to that observed without the RIS. Improvements in path gains are evident across all UEs when examining scenarios with 100 or 1000 RIS elements.

In Table I, we present a catalog of the various *network configurations* evaluated within the scope of this work. Configs. I-IV focus on a single UE, whereas configs. V-VIII address the multi-UE case study, where scheduling decisions influence intra-slice resource allocation. Regarding configs. VII-VIII, an xApp tasked with reconfiguring the scheduling profile selection is evaluated. The xApp updates the scheduling profile of the BS at a granularity of 1 s, assigning one of three scheduling profiles (i.e., RR, WF, and PF) to the UEs of the eMBB and URLLC slices, cycling through all possible combinations.

#### A. Single-UE Case Study A

In Fig. 3, we compare the UE with the best channel conditions, as shown in Fig. 2 (i.e., UE with ID 1), to the UE with the worst (i.e., UE with ID 5). In Fig. 3(a), focusing on Configs. I and II (Table I), which correspond to the UE with the worst channel, we observe that the median throughput reported with 100 RIS elements is 3.834 Mbps, representing an increase of  $\sim 10\%$  compared to the 3.486 Mbps observed with no RIS. Similarly, when focusing on the UE with the best channel (Configs. III and IV), we observe that the median throughput

<sup>1</sup>It is noted that all the results presented in Section VI have been averaged over multiple repetitions of the experiments.

TABLE	Ι
CATALOG OF NETWORK	CONFIGURATIONS

Config. ID	Number of UEs	UE ID(s)	Slice Type	Bandwidth [MHz]	RIS Assisted	xApp Enabled
I	1	5	eMBB	3.6	No RIS	No
Π	1	5	eMBB	3.6	100 RIS el.	No
Ш	1	1	eMBB	3.6	10 RIS el.	No
IV	1	1	eMBB	3.6	$1000~\mathrm{RIS}$ el.	No
V	5	$\{1,2\}$ & $\{3,4,5\}$	eMBB & URLLC	5 <b>&amp;</b> 5	No RIS	No
VI	5	$\{1,2\}$ & $\{3,4,5\}$	eMBB & URLLC	5 <b>&amp;</b> 5	100 RIS el.	No
VII	5	$\{1,2\}$ & $\{3,4,5\}$	eMBB & URLLC	9 <b>&amp;</b> 1	No RIS	Yes
VIII	5	$\{1,2\}$ & $\{3,4,5\}$	eMBB & URLLC	9 <b>&amp;</b> 1	100 RIS el.	Yes



Fig. 5. Performance evaluation results for the multi-UE O-RAN Case Study.



Fig. 6. Performance evaluation results for the multi-UE Case Study with and without the inclusion of the Sched xApp in O-RAN.

reported with 1000 RIS elements is 4.154 Mbps, representing an ~ 10% increase compared to the 3.8 Mbps achieved with only 10 RIS elements. It is reminded that according to the experimental results presented in Fig. 2, the performance achieved with 10 RIS elements is nearly identical to that observed without a RIS in the topology. Based on the aforementioned experimental evaluation, we observe that increases in the order of  $10^2$  in the size of the panels (e.g., from no RIS to  $10^2$  RIS, or from  $10^1$  RIS to  $10^3$  RIS) can result in a ~ 10% increase in the reported throughput under the same resource availability (e.g., 18 PRBs or 3.6 MHz bandwidth).

In Figs. 3(b) and 3(c), we present the corresponding Channel Quality Information (CQI) and Modulation and Coding Scheme (MCS) values as reported by the protocol stack. The results in Fig. 3(b) demonstrate the channel improvement introduced by the RIS in the topology, particularly for the UE with the best channel conditions (Configs. **III-IV**), which reports a  $\sim 13\%$  improvement in the median CQI value (from 8 to 9). Similarly, the MCS (Fig. 3(c)) increases from a median value of 13 to 16, enabling higher data rates given the same PRB availability. Regarding the UE with the worst channel (Configs. **I-II**), the MCS shows an increase of  $\sim 18\%$ , rising from 11 to 13, while the CQI remains relatively stable around 7, with a modest increase of  $\sim 8\%$  for the RIS topology.

#### B. Multi-UE Case Study B

We examine the multi-UE case study, where UEs are allocated to specific slices based on their traffic demands. To express UE satisfaction (which corresponds to high resource utilization), we define a new metric, i.e., the PRB Ratio. Its definition is given as PRB Ratio Sum of Granted PRBs Sum of Requested PRBs, where PRB Ratio  $\in [0, 1]$ , and it represents the amount of allocated PRBs to the slice according to the channel and traffic conditions. The analysis focuses on Configs. V and VI (Table I), where the UEs with the best reported channel conditions (Fig. 2) are grouped together on the eMBB slice, while the remaining UEs are assigned to the URLLC. The WF algorithm is selected for the scheduling profile of the eMBB (which has the highest throughput requirements), while RR is used for the URLLC (which has the lowest). Fig. 4 presents the complete performance evaluation results for the aforementioned case. Fig. 4(a) illustrates the median eMBB throughput for the RIS and non-RIS assisted topologies. The median throughput with a RIS is 2.973 Mbps,  $\sim 34\%$  higher than the 2.219 Mbps achieved without a RIS, under the same resource availability. Fig. 4(b) shows that both topologies (with and without a RIS) result in a median buffer occupancy of 0 in the URLLC, indicating significantly reduced latency levels.<sup>2</sup> Fig. 4(c) shows that although both topologies indicate near-zero latency, the inclusion of a RIS (Config. VI) leads to higher utilization efficiency of the available PRBs at the BS's scheduler.

# C. Multi-UE O-RAN Case Study C

In Fig. 5, focusing on Configs. VII and VIII, as detailed in Table I, we present the results of the experimental evaluation where the Sched xApps controls the BS's scheduling policy. Fig. 5(a) shows that the median eMBB throughput with a RIS in the topology is 4.234 Mbps, compared to 4.1 Mbps without a RIS. These high throughput values are reasonable, as the eMBB UEs are allocated 45 PRBs, representing 90%of the available resources. Fig. 5(b), demonstrates that in the presence of a RIS, the median URLLC buffer occupancy is 0, whereas in its absence, the corresponding value is 688, indicating high latency levels. Therefore, the results show that the presence of 100 RIS elements significantly reduces latency levels, regardless of the BS's scheduling policy. On the contrary, non-RIS-assisted topologies in resource-limited environments (i.e., with only  $\sim 10\%$  of the available bandwidth) experience performance degradation, even when the BS frequently updates its scheduler from fair (e.g., RR) to more efficient approaches (e.g., WF or PF) to accommodate the UEs based on their distinct Quality of Service (QoS) demands. Finally, Figs. 5(c) and 5(d) demonstrate that RIS-enabled topologies achieve higher resource utilization for both the eMBB and URLLC slices. Figs. 6(a) and 6(b) compare the previously discussed topologies to the case where no xApp reconfigures the BS's scheduling policy. Notably, the keyword baseline, as shown in Fig. 6, refers to srsRAN's default scheduling policy, which is RR. Fig. 6(a) shows that updating the BS's scheduling profile increases resource efficiency for URLLC in the absence of a RIS (Config. VII). However, in the presence of a RIS (Fig. 6(b)), resource utilization and hence UE satisfaction improve significantly (Config. VIII). Therefore, both the presence of RIS in the topology and frequent updates of the BS's scheduling policy enhance the resource utilization efficiency.

# VII. CONCLUSIONS AND FUTURE WORK

In this work, we focused on an Open RAN and investigated the impact of various network management policies on the performance of RIS-assisted topologies, leveraging the capabilities of Open RAN Gym on Colosseum. Our experimental results demonstrate that a RIS deployment can enhance network throughput and resource utilization under the same bandwidth availability. In addition, our experimentation in an emulated O-RAN system, where the BS's scheduling policy is controlled by an xApp, demonstrates that RIS-assisted topologies achieve superior network performance and higher resource efficiency compared to their non-RIS-assisted counterparts, regardless of the scheduling policy employed. Part of our current and future work focuses on designing AI/ML solutions to dynamically control the selection of scheduling policies in such Open RANs, leveraging CSI and traffic conditions in the decision-making processes. Extensions to the O-RAN control plane for carrying RIS weights, determined by the scheduler for real-time network adaptation, are also left for future work.

#### REFERENCES

- 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*, 2023.
- [2] AI-RAN Alliance, "AI-RAN Alliance Vision and Mission White Paper," [Online]. Available: https://ai-ran.org/wp-content/uploads/2024/12/ AI-RAN\_Alliance\_Whitepaper.pdf, December 2024.
- [3] C. Liaskos, L. Mamatas, A. Pourdamghani, A. Tsioliaridou, S. Ioannidis, A. Pitsillides, S. Schmid, and I. F. Akyildiz, "Software-defined reconfigurable intelligent surfaces: From theory to end-to-end implementation," *Proceedings of the IEEE*, vol. 110, no. 9, pp. 1466–1493, 2022.
- [4] J. Angjo, A. Zubow, and F. Dressler, "Side Effects of IRS: On the Need for Coordination in 6G Multi-Operator IRS-assisted Networks," in *IEEE Global Communications Conference (GLOBECOM), 4th Workshop on Emerging Topics in 6G Communications (6GComm).* IEEE, 12 2023, pp. 1380–1385.
- [5] L. Bonati, M. Polese, S. D'Oro, S. Basagni, and T. Melodia, "OpenRAN Gym: AI/ML development, data collection, and testing for O-RAN on PAWR platforms," *Computer Networks*, vol. 220, p. 109502, 2023.
- [6] M. Tsampazi, S. D'Oro, M. Polese, L. Bonati, G. Poitau, M. Healy, M. Alavirad, and T. Melodia, "PandORA: Automated Design and Comprehensive Evaluation of Deep Reinforcement Learning Agents for Open RAN," *IEEE Transactions on Mobile Computing*, 2024.
- [7] S. Li, B. Duo, X. Yuan, Y.-C. Liang, and M. Di Renzo, "Reconfigurable intelligent surface assisted UAV communication: Joint trajectory design and passive beamforming," *Wireless Communications Letters*, vol. 9, no. 5, pp. 716–720, 2020.
- [8] M. Diamanti, M. Tsampazi, E. E. Tsiropoulou, and S. Papavassiliou, "Energy efficient multi-user communications aided by reconfigurable intelligent surfaces and UAVs," in *IEEE International Conference on Smart Computing (SMARTCOMP)*, 2021, pp. 371–376.
- [9] M. Diamanti, P. Charatsaris, E. E. Tsiropoulou, and S. Papavassiliou, "The prospect of reconfigurable intelligent surfaces in integrated access and backhaul networks," *IEEE Transactions on Green Communications and Networking*, vol. 6, no. 2, pp. 859–872, 2021.
- [10] M. Tsampazi and T. Melodia, "System-Level Experimental Evaluation of Reconfigurable Intelligent Surfaces for NextG Communication Systems," arXiv preprint arXiv:2412.12969, 2024.
- [11] M. Polese, L. Bonati, S. D'Oro, P. Johari, D. Villa, S. Velumani, R. Gangula, M. Tsampazi, C. P. Robinson, G. Gemmi *et al.*, "Colosseum: The open RAN digital twin," *IEEE Open Journal of the Communications Society*, 2024.
- [12] ETSI, "Reconfigurable Intelligent Surfaces (RIS); Use Cases, Deployment Scenarios and Requirements," ETSI, ETSI Group Report RIS 001 V1.1.1, April 2023. [Online]. Available: https://www.etsi.org/ deliver/etsi\_gr/RIS/001\_099/001/01.01.01\_60/gr\_RIS001v010101p.pdf
- [13] S. Jaeckel, L. Raschkowski, K. Börner, and L. Thiele, "QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials," *IEEE transactions on antennas and propagation*, vol. 62, no. 6, pp. 3242–3256, 2014.
- [14] B. Zheng, Q. Wu, and R. Zhang, "Intelligent Reflecting Surface-Assisted Multiple Access With User Pairing: NOMA or OMA?" *IEEE Communications Letters*, vol. 24, no. 4, pp. 753–757, 2020.
- [15] I. Gomez-Miguelez, A. Garcia-Saavedra, P. D. Sutton, P. Serrano, C. Cano, and D. J. Leith, "srsLTE: An Open-Source Platform for LTE Evolution and Experimentation," in *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization*, 2016, pp. 25–32.
- [16] U.S. Naval Research Laboratory, "Multi-Generator (MGEN) Network Test Tool". https://www.nrl.navy.mil/itd/ncs/products/mgen. 2019.

<sup>&</sup>lt;sup>2</sup>It is noted that since srsRAN does not directly measure latency, buffer occupancy is used as a proxy [6].