

# OpenRAN Gym: AI/ML Development, Data Collection, and Testing for O-RAN on PAWR Platforms

Leonardo Bonati, Michele Polese, Salvatore D’Oro, Stefano Basagni, Tommaso Melodia

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

E-mail: {bonati.l, m.polese, s.doro, s.basagni, melodia}@northeastern.edu

**Abstract**—Open Radio Access Network (RAN) architectures will enable interoperability, openness and programmable data-driven control in next generation cellular networks. However, developing and testing efficient solutions that generalize across heterogeneous cellular deployments and scales, and that optimize network performance in such diverse environments is a complex task that is still largely unexplored. In this paper we present OpenRAN Gym, a unified, open, and O-RAN-compliant experimental toolbox for data collection, design, prototyping and testing of end-to-end data-driven control solutions for next generation Open RAN systems. OpenRAN Gym extends and combines into a unique solution several software frameworks for data collection of RAN statistics and RAN control, and a lightweight O-RAN near-real-time RAN Intelligent Controller (RIC) tailored to run on experimental wireless platforms. We first provide an overview of the various architectural components of OpenRAN Gym and describe how it is used to collect data and design, train and test artificial intelligence and machine learning O-RAN-compliant applications (xApps) at scale. We then describe in detail how to test the developed xApps on softwarized RANs and provide an example of two xApps developed with OpenRAN Gym that are used to control a network with 7 base stations and 42 users deployed on the Colosseum testbed. Finally, we show how solutions developed with OpenRAN Gym on Colosseum can be exported to real-world, heterogeneous wireless platforms, such as the Arena testbed and the POWDER and COSMOS platforms of the PAWR program. OpenRAN Gym and its software components are open-source and publicly-available to the research community.

## I. INTRODUCTION

Once seen as monolithic and mostly immutable “black-box” systems, cellular networks are converging toward the more flexible, software-based open architectures based on the Open Radio Access Network (RAN) paradigm. This new approach to cellular communications promotes openness, virtualization, and programmability of RAN functionalities and components, and enables data-driven intelligent control loops for cellular systems [2]. As such, the Open RAN enables network operators to support new bespoke services on shared physical infrastructures, and to dynamically reconfigure them based on network conditions and user demand. The resulting

This is a revised and substantially extended version of [1], which appeared in the Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC) 2022 Workshops.

This work was partially supported by the U.S. National Science Foundation under Grants CNS-1925601, CNS-2120447, and CNS-2112471.

increased efficiency will also decrease the operational costs of the network.

In this context, standardization bodies and other organizations are releasing a number of specifications to regulate the operations of the Open RAN, and to define its capabilities, constraints, and use cases. The most notable is the O-RAN Alliance, which is developing specifications—collected under the O-RAN umbrella—to apply Open RAN principles to prevailing radio access technologies, including 3GPP LTE and NR networks [3].

O-RAN introduces two network RAN Intelligent Controllers (RICs), operating at different timescales, enabling programmatic closed-loop control of the RAN elements. It also defines a set of open interfaces to connect the controllers to key elements of the RAN, such as the NR Central Units (CUs), Distributed Units (DUs), Radio Units (RUs), and the LTE O-RAN-compliant evolved Node Bases (eNBs) [4]. In details, the near-real-time (or near-RT) RIC connects to the RAN elements (i.e., the CUs and DUs) through the E2 interface, and enables control loops operating at timescales ranging between 10 ms and 1 s [5]. Instead, the non-real-time (or non-RT) RIC is included as part of Service Management and Orchestration (SMO) frameworks, and operates at timescales larger than 1 s [6]. This component also interacts with one or multiple the near-RT RICs via the A1 interface, which is used to disseminate policies and information external to the network. The non-RT RIC also manages the Artificial Intelligence (AI) and Machine Learning (ML) models which are instantiated on the RICs in the form of standalone applications, namely xApps (on the near-RT RIC) and rApps (on the non-RT RIC). Finally, the SMO connects to the RAN through the O1 interface, used for management and orchestration routines, and to the O-RAN virtualization platform (the O-Cloud) via the O2 interface.

Thanks to its RICs, open interfaces and disaggregated architecture, O-RAN ultimately enables the practical deployment and execution of AI/ML solutions at scale, which can be used to infer and forecast network traffic, or to reconfigure the nodes of the RAN at run time based on real-time conditions and user demand. Typical workflows for the design and testing of such AI/ML algorithms encompass a number of different steps such as [7, 8]): (i) *data collection*, to create practical datasets representative of the different environments (e.g., the wireless channel) where the AI/ML models will be deployed, as well as of various performance indicators of the network;

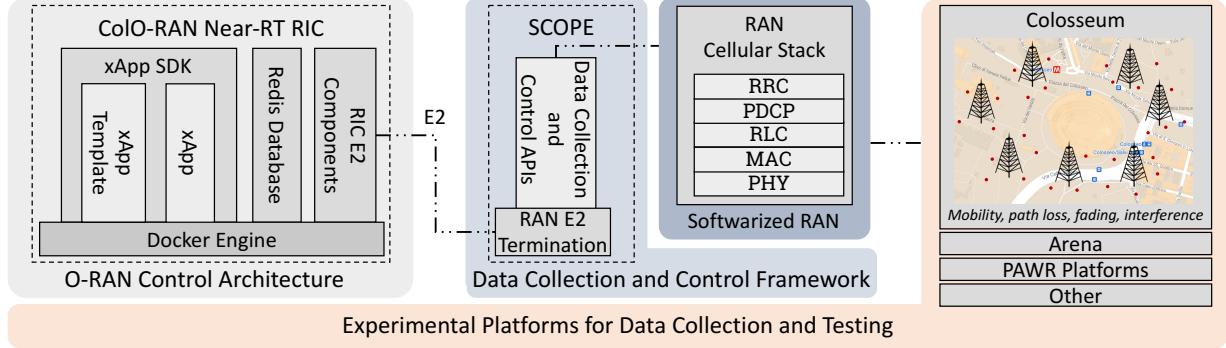


Figure 1: OpenRAN Gym architecture.

(ii) *AI/ML model design*, selecting the inputs and outputs of the models, and *training and testing*, to evaluate the effectiveness and limits of such models; (iii) *model deployment* as applications deployed on the RICs, i.e., xApps/rApps or—as recently proposed in [9]—directly on the CUs/DUs via dApps; (iv) *model fine-tuning* with run-time data from the RAN, to adapt the models to different production environments, and (v) the actual *control, inference and/or forecasting* of the RAN.

In this paper we present OpenRAN Gym, an open-source toolbox to develop AI/ML O-RAN-compliant inference and control algorithms, to deploy them as xApps on the near-RT RIC, and to test them on a large-scale softwarized RAN controlled by the RIC. First, we give a high-level overview of the various components of OpenRAN Gym, and discuss how they enable development and testing workflows of data-driven xApps. We showcase an example of two xApps designed with OpenRAN Gym and used to control a large-scale RAN instantiated on the Colosseum wireless network emulator [10] through the SCOPE framework [11], and controlled by the Colo-RAN near-RT RIC [8]. We also show how OpenRAN Gym can be seamlessly ported from an emulator such as Colosseum to over-the-air real-world platforms, such as the Arena testbed [12], and the platforms of the U.S. National Science Foundation-sponsored Platforms for Advanced Wireless Research (PAWR) program [13] including the Platform for Open Wireless Data-driven Experimental Research (POW-DER) [14] and the Cloud Enhanced Open Software Defined Mobile Wireless Testbed for City-Scale Deployment (COSMOS) [15] platforms. To the best of our knowledge, OpenRAN Gym is the first open, portable toolset for end-to-end design, prototyping, testing, and experimentation of AI/ML O-RAN xApps on heterogeneous wireless experimental platforms.

Previous experimental work has focused on the development of data-driven solutions and xApps for specific use cases [16, 17], on the description of the AI/ML capabilities of O-RAN [18, 19], on interoperability testing [20], and on orchestration [21]. Compared to the state of the art, OpenRAN Gym enables an end-to-end workflow for the design and testing of AI/ML solutions as xApps in the O-RAN ecosystem. By doing so, it empowers users with a first-of-its-kind open and publicly-available O-RAN-compliant toolbox that will unleash the potential of data-driven applications for next generation cellular networks. Being open-source, OpenRAN Gym aims at

creating a thriving community of researchers and developers contributing to it with open-source software components for experimental O-RAN-enabled data-driven research.<sup>1</sup>

The remainder of this paper is organized as follows. We give an overview of the various components of OpenRAN Gym in Section II. Practical descriptions of OpenRAN Gym data collection and control framework, and O-RAN control architecture, are given in Sections III and IV, respectively. The xApp design and testing workflow is presented in Section V, along with an example of large-scale RAN control using xApps developed with OpenRAN Gym on Colosseum. Section VI discusses how OpenRAN Gym components and experiments can be ported from Colosseum to heterogeneous real-world testbeds. Section VII showcases exemplary results obtained on the different platforms considered in this work. Finally, conclusions are drawn in Section VIII.

## II. OPENRAN GYM

The OpenRAN Gym architecture is shown in Figure 1. Its main components are: (i) publicly- and remotely-accessible *experimental wireless platforms* for collecting data, prototyping, and testing solutions in heterogeneous environments. Examples of these are the Colosseum wireless network emulator [10], the Arena testbed [12], and the platforms of the PAWR program [13]; (ii) a *softwarized RAN* implemented through open protocol stacks for cellular networks, such as srsRAN [22] and OpenAirInterface [23]; (iii) a *data collection and control framework*, such as SCOPE [11], that exposes Application Programming Interfaces (APIs) to extract relevant Key Performance Measurements (KPMs) from the RAN, and dynamically control it at run-time, and (iv) an *O-RAN control architecture*, such as Colo-RAN [8], able to connect to the RAN through open and standardized interfaces (e.g., the O-RAN E2 interface), receive the run-time KPMs from the RAN, and control it through AI/ML solutions running, for instance, as xApps/rApps. As we will show in Sections VI and VII, OpenRAN Gym is platform-independent, and it allows users to perform data collection campaigns, prototype, and evaluate solutions in a set of heterogeneous wireless environments and deployments before transitioning them to production networks. As such, OpenRAN Gym can be used to first prototype and

<sup>1</sup>The software components of OpenRAN Gym are publicly-available and accessible at <https://openrangym.com>.

validate solutions on the Colosseum wireless network emulator, and then seamlessly transfer such solutions to heterogeneous platforms, such as the Arena testbed, and the POWDER and COSMOS platforms from the PAWR program [13]. The procedures to port the various components of OpenRAN Gym on these platforms will be described in Section VI.

Arena is an indoor wireless testbed equipped with a grid of 64 antennas and 24 Software-defined Radios (SDRs) (among USRPs X310 and N210) controlled by high-performance compute servers [12]. Its deployment is representative of a live office environment.

Colosseum is the world’s largest wireless network emulator [10]. It allows researchers and practitioners to experiment at scale, and in different channel conditions and virtual environments through a set of 128 SDRs (USRPs X310) controlled through dedicated servers—namely, Standard Radio Nodes (SRNs)—interconnected through a Massive Channel Emulator (MCHEM). The latter, is capable of reproducing conditions of the wireless channel (e.g., path loss, fading, user mobility, signal interference and superimposition) by means of Finite Impulse Response (FIR) filters implemented through Field Programmable Gate Arrays (FPGAs). The channel emulation is performed by the FIR filters, which apply the channel impulse response of the desired wireless channel to the signals transmitted by the SRNs. Sets of channel impulse responses for different environments (e.g., urban, rural, etc.)—referred to as *Radio Frequency (RF) scenarios* in Colosseum—are modeled a priori through mathematical equations, or captured through ray-tracing software.

The POWDER is a city-scale wireless testbed deployed in Salt Lake City, UT [14]. The testbed includes a number of SDRs deployed across an outdoor area, an over-the-air indoor laboratory setup, and a wired attenuator matrix. The objective of this testbed is to foster experimental research in heterogeneous technology, such as 5G cellular technologies and network orchestration.

The COSMOS is a city-scale testbed deployed in New York City, NY, which mainly focuses on mmWave communications with edge-computing capabilities [15]. This testbed absorbed the Open-Access Research Testbed for Next-Generation Wireless Networks (ORBIT) [24], an indoor over-the-air wireless platform with remotely-accessible SDR devices and compute servers.

At the time of this writing, OpenRAN Gym softwarized RAN leverages the cellular implementation provided by srsRAN [22], which allows users to instantiate protocol stacks of 3GPP base stations and User Equipments (UEs) using SDRs as front-end interfaces. This cellular protocol stack is augmented by the SCOPE framework, which adds a number of networking and control functionalities to srsRAN including network slicing capabilities, support for additional scheduling algorithms, data collection pipelines, and open APIs to control such functionalities at run time. As we will discuss in Section III, SCOPE can facilitate data collection campaigns by automating the collection of relevant RAN KPM in the heterogeneous testbed where it is instantiated [8, 21, 25].

Finally, Colo-RAN implements the O-RAN control architecture of OpenRAN Gym. This framework adapts the near-RT

RIC provided by the O-RAN Software Community (OSC) to run in a lightweight containerized environment, and extends it to swiftly interface with, and control, the SCOPE base stations through the E2 interface standardized by O-RAN. As we will discuss in Section IV, Colo-RAN allows users to prototype AI/ML-based O-RAN applications through an *xApp Software Development Kit (SDK)*, to instantiate them on an OSC-compliant near-RT RIC, and to leverage them to perform control of a softwarized RAN (Figure 1).

### III. DATA COLLECTION AND CONTROL FRAMEWORK

The data collection and control framework of OpenRAN Gym is based on SCOPE [11]. This framework provides a programmable environment for prototyping and testing solutions for softwarized RANs, and data collection capabilities of relevant KPMs (e.g., throughput, Transport Blocks (TBs), buffer occupancy). Concerning the cellular protocol stack for base stations and UEs, SCOPE leverages srsRAN [22], which it extends with novel network slicing and a set of additional scheduling algorithms. Open APIs to fine-tune the configuration of the RAN at run time, and to perform data collection campaigns are also provided by SCOPE. Coupled with different testbeds—such as Colosseum and the platform of the PAWR program—SCOPE can facilitate the collection of RAN KPMs in a set of heterogeneous scenarios and environments by automatically collect such statistics from the running experiments [8, 11, 25]. Finally, SCOPE connects to the O-RAN near-RT RIC through a RAN-side O-RAN E2 termination, which is based on the OSC DU [26]. This allows user-defined xApps running on the near-RT RIC to swiftly interface with the RAN base stations, and to dynamically control their functionalities at run time (e.g., modify the scheduling policy and set the amount of resources allocated to each network slice). In the remainder of this section, we will give a high-level overview of the main configuration options and parameters of the SCOPE-enabled base stations, and show how to instantiate a cellular network with it. SCOPE has been open-sourced to the research community,<sup>2</sup> and also provided to the Colosseum users in the form of a ready-to-use Linux Container (LXC) (namely `scope`/`scope-with-e2`). In Section VI, we will show how the publicly-available SCOPE container can be ported to different testbeds (e.g., the Arena testbed, and the POWDER and COSMOS testbeds of the PAWR program) with minor modifications. In this way, SCOPE truly enables the process of cellular-network-as-a-service, in which the solutions are first prototyped in a controlled environment (e.g., Colosseum), and then ported in the wild on real-world testbeds.

#### A. Starting SCOPE

SCOPE provides Command-line Interface (CLI) tools to start the cellular base stations and configure them through parameters passed via configuration files. The main parameters of interest to OpenRAN Gym are described as follows.<sup>3</sup>

<sup>2</sup>The SCOPE source code is available at <https://github.com/wineslab/colosseum-scope> and <https://github.com/wineslab/colosseum-scope-e2>.

<sup>3</sup>A comprehensive description of the SCOPE APIs and configuration parameters can be found at <https://github.com/wineslab/colosseum-scope>.

- **network-slicing**: enables/disables the network slicing functionalities of the base station.
- **slice-allocation**: if network slicing has been enabled, this parameter can be used to set the Resource Block Groups (RBGs) allocated by the base station to each slice. The input of this configuration option is passed as `{slice:[first_rbg, last_rbg],...}`. As an example, `{0:[0,5],1:[6,10]}` allocates RBGs 0-5 to slice 0 and 6-10 to slice 1.
- **slice-scheduling-policy**: sets the scheduling policy used for each network slice of the base station. As an example, `[1,2]` assigns slicing policy 1 to slice 0 and policy 2 to slice 1. The possible numerical values for this field match the scheduling policies supported by SCOPE (i.e., 0: round-robin, 1: waterfilling, 2: proportionally fair).
- **slice-users**: associates UEs to a specific network slice. The input of this configuration option is passed as `{slice:[ue1,ue2],...}`. As an example, `{0:[4,5],1:[2,3]}` assigns UEs 4, 5 to slice 0, and UEs 2, 3 to slice 1.
- **generic-testbed**: specifies whether SCOPE is running on a testbed other than Colosseum. In this case, the parameters `node-is-bs` and `ue-id` can also be passed to specify whether the node should act as a base station or a UE, and the identifier of the UE in the latter case.<sup>4</sup>

After the SCOPE configuration has been written in a JSON-formatted file,<sup>5</sup> (named `radio.conf` in the code snippet below), the cellular base station, core network, and UE applications can be started through the commands of Listing 1.

---

```
1 #!/bin/bash
2 cd radio_api/
3 python3 scope_start.py --config-file radio.conf
```

---

Listing 1: Commands to start the SCOPE applications.

At run-time, the SCOPE APIs can be leveraged to fine-tune the configuration of the base station, e.g., to modify the scheduling policy of each slice, or to set the amount of RBGs of each slice (see [11, Section 3.3]). Relevant KPMs from the RAN are automatically logged by the SCOPE base stations while traffic is exchanged among base stations and UEs. These KPMs are saved in CSV-formatted files that can be either used on-the-fly (e.g., for online AI/ML model training or inference), or retrieved for offline processing after the experiment ends (e.g., to perform offline AI/ML model training).

#### IV. O-RAN CONTROL ARCHITECTURE

The O-RAN control architecture leveraged by OpenRAN Gym is based on Colo-RAN, an open-source framework to develop, design, prototype, and test O-RAN-ready solutions at scale [8]. This framework provides a lightweight implementation of the OSC near-RT RIC—which has been adapted

<sup>4</sup>When running on Colosseum, SCOPE automatically derives the role of the node, and the UE identifier based on the allocated SRNs.

<sup>5</sup>An example of configuration file can be found at <https://tinyurl.com/2s3pvw83> (for Colosseum), and at <https://tinyurl.com/35t7s97a> (for testbeds other than Colosseum).

to run on the Colosseum system as a set of standalone Docker containers—as well as automated pipelines for the deployment of the various services of the RIC. The main components implemented by Colo-RAN are shown in Figure 1, left. They are services in charge of overseeing the interactions with the RAN (e.g., the *E2 termination*, *E2 manager*, and *E2 routing manager*), a *Redis database* that keeps records of the connected RAN nodes (e.g., the base stations), and an *xApp SDK* with tools to prototype and test AI/ML-based xApp for run-time RAN inference and/or control.

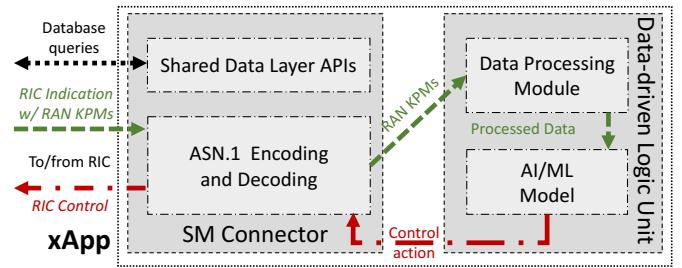


Figure 2: Colo-RAN xApp, adapted from [8].

A high-level diagram of a Colo-RAN xApp is shown in Figure 2. This is formed of two main parts: (i) the *Service Model (SM) connector*, in charge of handling the messages between the xApp and the near-RT RIC (e.g., the control messages for the RAN), and (ii) the *data-driven logic unit*, which processes KPMs received from the RAN base stations, and performs tasks based on AI/ML models (e.g., traffic prediction and/or control of the functionalities of the base stations). The data-driven logic unit hosts two sub-components, namely the *AI/ML model* and the *data processing module*. The former consists of the specific data-driven model (e.g., Deep Reinforcement Learning (DRL) agent, Deep Neural Network (DNN), Long Short Term Memory (LSTM), to name a few) used to perform inference and/or control tasks. The latter, instead, executes data processing functionalities to convert the input KPMs into data that can be fed to the AI/ML model. For instance, the majority of AI/ML models are designed to receive inputs with a fixed size and format (e.g., a two-dimensional array of a specific length, an image, or a time-series). However, the KPMs received over the E2 termination might have a different format, or might contain more data than what is required by the AI/ML model. In this case, the data processing module performs the necessary operations to convert the input data in the correct format. In some cases, the data processing module can also host some AI/ML models that execute advanced data processing operations. Examples of these are autoencoders to extract latent data representation and to perform dimensionality reduction [8, 25].

In the remainder of this section, we detail how to instantiate the Colo-RAN near-RT RIC (Section IV-A), how to interface it with the SCOPE base station through the O-RAN E2 termination (Section IV-B), and how to start a sample xApp that controls the base station (Section IV-C). Colo-RAN has been open-sourced and made available to the

research community,<sup>6</sup> and also provided to the Colosseum users in the form of a ready-to-use LXC container (namely `coloran-near-rt-ric`). In Section VI, we will show how ColO-RAN can be seamlessly ported to different testbeds (e.g., the Arena testbed, and the POWDER and COSMOS platforms of the PAWR program).

#### A. Starting the ColO-RAN Near-RT RIC

The ColO-RAN near-RT RIC can be built and instantiated as a set of Docker containers by running the `setup-ric.sh` script and the commands of Listing 2.<sup>7</sup> This script, adapted from [27], takes as input the network interface the RIC uses to receive and exchange messages with the RAN (e.g., the `col0` interface in Colosseum). As a first step, the base Docker

---

```

1 #!/bin/bash
2 cd setup-scripts/
3 ./setup-ric.sh col0

```

---

Listing 2: Commands to set up the ColO-RAN near-RT RIC.

images that will be used to build the RIC are imported. Then, the actual Docker images of the ColO-RAN near-RT RIC are build. These images include: (i) the `e2term`, which is the endpoint of the RIC E2 messages; (ii) the `e2mgr`, which is in charge of managing the messages to/from the E2 interface; (iii) the `e2rtmansim`, which uses the RIC Message Router (RMR) protocol to route the E2 messages within the RIC; and (iv) the `db`, which implements a Redis database with records of the RAN nodes connected to the RIC (e.g., the base stations). During this step, the IP addresses and ports that will be used by the Docker containers are also configured as set up in the `setup-lib.sh` file. After the Docker images have been built, the RIC containers, listening for incoming connections from the RAN through the E2 termination endpoint, are spawned. The logs of the various containers can be accessed through the `docker logs` command, e.g., `docker logs -f e2term` shows the run-time logs of the E2 termination (`e2term`) container.

#### B. Connecting the SCOPE Base Station to ColO-RAN

After setting up and starting ColO-RAN through the steps described in Section IV-A, the cellular base station—provided by SCOPE and set up in Section III-A—can be connected to it through the O-RAN E2 termination, which has been adapted from the OSC DU implementation [26]. To this aim, the RAN-side E2 termination can be used to: (i) receive *RIC Subscription* messages from the xApps; (ii) transmit periodic KPM reports to the xApps through *RIC Indication* messages, and receive control actions from them through *RIC Control* messages, and (iii) interact with the APIs provided by SCOPE to modify the configuration of the base station at run time (e.g., the scheduling and slicing policies) based on the control messages received from the xApps. The steps to initialize

<sup>6</sup>The ColO-RAN source code is available at <https://github.com/wineslab/colosseum-near-rt-ric>.

<sup>7</sup>A pre-built version of ColO-RAN, named `coloran-near-rt-ric-prebuilt`, is also provided to the Colosseum users. The pre-built Docker images are also hosted on Docker Hub at <https://hub.docker.com/u/wineslab>.

the E2 termination at the SCOPE base station are shown in Listing 3. First, the E2 termination is build through the

---

```

1 #!/bin/bash
2 cd colosseum-scope-e2/
3 ./build_odu.sh clean
4 ./run_odu.sh

```

---

Listing 3: Commands initialize the SCOPE E2 termination process.

`build_odu.sh` script (line 3). This script also specifies the IP address and port of the near-RT RIC to connect to, as well as the local network interface used for the connection. Then, the E2 termination process is started through the `run_odu.sh` script (line 4), which establishes the initial connection between the base station and the near-RT RIC. The successful outcome of this connection can be verified in the logs of the `e2term` container (via the `docker logs -f e2term` command, see Section IV-A), which reports the identifier of the connected base station (e.g., `gnb:311-048-01000501`).

#### C. Initializing a Sample xApp

After the SCOPE base station has been connected to the near-RT RIC, the xApps can be started. To facilitate the design of novel xApps, we provide a ready-to-use sample xApp template in which researchers and practitioners can plug-in their custom AI/ML models. This sample xApp can be started through the `setup-sample-xapp.sh` script and the commands shown in Listing 4. This script takes as input the

---

```

1 #!/bin/bash
2 cd setup-scripts/
3 ./setup-sample-xapp.sh gnb:311-048-01000501

```

---

Listing 4: Commands to build the ColO-RAN sample xApp Docker image, and to start and configure the xApp container.

identifier of the base station the xApp should subscribe to (see Section IV-B), and builds the Docker image of the sample xApp. Then, the script starts the xApp Docker container—dubbed `sample-xapp`—on the near-RT RIC.

After the container has started, the xApp processes can be run through the commands of Listing 5. These commands

---

```

1 #!/bin/bash
2 docker exec -it sample-xapp
   /home/sample-xapp/run_xapp.sh

```

---

Listing 5: Commands to run the ColO-RAN sample xApp process.

trigger the xApp subscription to the targeted RAN nodes (e.g., one or multiple base stations connected to the RIC) through RIC Subscription messages, and the periodic reports of RAN KPMs from such nodes. Starting from the provided template, OpenRAN Gym users can build xApps running custom solutions (e.g., with custom AI/ML agents).

## V. xAPP DEVELOPMENT WORKFLOW ON COLOSSEUM

The main steps to develop a data-driven xApp using OpenRAN Gym on Colosseum are shown in Figure 3.

**1) Data collection.** This step involves collecting the data that will be used to train and test the AI/ML model to embed in the xApp. In Colosseum, this can be done by combining

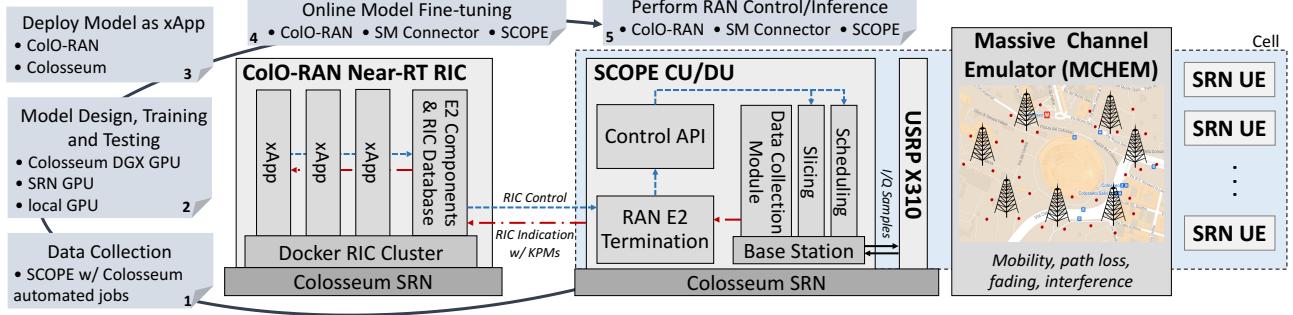


Figure 3: OpenRAN Gym xApp design and testing workflow on Colosseum.

the data-collection capabilities of SCOPE with the automated experiments of Colosseum. This allows to automatically run experiments with several base stations and users in a set of heterogeneous scenarios, and to collect the RAN KPMs—saved in CSV-formatted files—from Colosseum Network Attached Storage (NAS) once the experiment ends [11].

**2) Model design, training and testing.** After data collection campaigns have been performed in heterogeneous wireless environments and scenarios, the AI/ML model can be designed. This step includes the selection of the AI/ML algorithm that the model will use, along with the data used as input, the reward function, and the set of output actions (e.g., to perform inference or control of the RAN). After this design phase, the model is first trained offline using the data collected in step 1, and then tested at scale.<sup>8</sup> Being computationally-intensive, this step may benefit from Graphics Processing Unit (GPU)-enabled environments. As such, they can be carried out either locally (i.e., on the user’s own GPU-enabled machines), or on Colosseum’s GPU-enabled SRNs or NVIDIA A100 DGXs.

**3) Deploy the model as an xApp.** After the model has been tested (step 2), it can be deployed as an xApp on the ColO-RAN near-RT RIC by following the procedures of Section IV-C. Specifically, the AI/ML model is included in the *data-driven logic unit* of Colo-RAN xApp (see Figure 2) by modifying the provided xApp template. Finally, the modified xApp is build and instantiated on the near-RT RIC through the commands of Listings 4 and 5.

**4) Online model fine-tuning.** At run-time, the xApp communicates with the SCOPE base station through the near-RT RIC and the E2 termination. To this aim, the xApp first subscribes to the base station by sending it a RIC Subscription message. Then, it triggers periodic KPMs reports—with periodicity tunable based on the needs of the users [8]—from the base station. These reports are sent through RIC Indication messages, and they may be used by the xApp to fine-tune the model online, allowing it to adapt to varying wireless conditions and traffic demand. Once the model has been fine-tuned online, the Docker image of the xApp can be updated with the trained weights.

**5) Perform RAN control/inference.** At this stage, the xApp can be used in the a live infrastructure to perform inference

and/or control of the RAN. This entails the xApp transmitting the actions computed by the model to the SCOPE base station through RIC Control messages. Example of these are actions to modify the parameters and configuration of the base station, e.g., to modify the resources allocated to the slices of the network, or their scheduling policies. At the base station, these RIC Control messages—received through the O-RAN E2 interface—trigger the SCOPE control APIs of Figure 3, which apply the new policies to the configuration of the base station at run time. At this point, the xApps can be tested and validated on Colosseum. In Sections VI and VII, we will show how these newly developed xApps can be ported and instantiated on external wireless testbeds.

#### A. Example of xApps

For the sake of completeness, we now provide an example of two xApps designed, trained and tested with OpenRAN Gym on Colosseum. These xApps are used to control a cellular network with 7 base stations and 42 UEs instantiated on the Colosseum network emulator. Each base station is implemented through SCOPE and serves 6 UEs with different traffic requirements. The UEs are divided into two classes of traffic, allocated to different slices of the network: time-sensitive (e.g, Ultra Reliable and Low Latency Communications (URLLC)) and broadband (e.g., Enhanced Mobile Broadband (eMBB) and Machine-type Communications (MTC)).

The xApps implement DRL agents—trained on a dataset with 3.4 GB of RAN traces (and more than 73 hours of experiments) collected on Colosseum—that make control decisions on the configuration of the base station based on the received RAN KPMs (see [8]). The DRL agents considered in this paper implement a Proximal Policy Optimization (PPO) architecture that leverages an actor-critic structure. The actor network is trained to take actions according to the current state of the system, while the critic network is used during the training phase to evaluate the reward obtained by selecting a specific action in a certain state. Then, the critic network instructs the actor network on how valuable the action was, in this way steering the actor network toward actions that bring the highest reward for each state.

To showcase the impact of different design choices on the overall performance of the network, we trained two xApps with different action spaces. One xApp (named *sched*) controls the scheduling policies that a base station uses for specific classes of traffic. Another xApp (*sched-slicing*) is

<sup>8</sup>It is worth mentioning that the O-RAN specifications forbid the deployment of models that have not been trained offline beforehand. This is to shield the RAN from poor performance or outages [7].

instead operating over a larger action space as it controls both scheduling policies, and the resource allocated to each slice (i.e., the number of RBGs assigned to each class of traffic). In this example, both xApps aim at maximizing the amount of transmitted data belonging to the broadband traffic class (in this case measured by the number of downlink TBs transmitted successfully by the base station to the UEs), and to minimize the time packets belonging to the time-sensitive traffic class spend in the downlink buffer queues of the base station. As the protocol stack of the base station does not have a direct measurement of the end-to-end system latency, we use this buffer occupancy metric as our proxy for latency.

Figure 4 shows the Cumulative Distribution Function (CDF) of some RAN metrics measured at the base stations when the two xApps are instantiated on the Colo-RAN near-RT RIC and used to control the RAN. Specifically, Figure 4a shows the transmitted TBs for the broadband slice, while Figure 4b displays the downlink buffer occupancy of the time-sensitive slice. By acting on a large action set (i.e., the slice resource

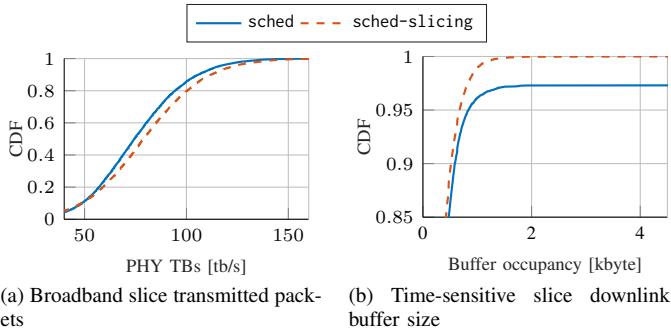


Figure 4: Comparison of xApps developed with OpenRAN Gym.

allocation), the *sched-slicing* xApp achieves superior performance by delivering a higher number of transmitted packets and reducing the occupancy of the downlink buffer.<sup>9</sup>

## VI. TRAVELING CONTAINERS

In this section, we illustrate how the OpenRAN Gym containerized applications including the xApps developed and pre-trained on Colosseum can be transferred to other testbeds, and describe the necessary adjustments (if any) to run these applications on each experimental platform. Although the above procedure may seem trivial in the case of self-contained applications, in our case this is challenging due to the fact that our *traveling OpenRAN Gym containers* need to interact with the underlying network resources and be able to properly control the potentially diverse set of SDRs available in the different testbeds. Furthermore, in some cases, the firmware of the SDRs requires specific tools only available on certain operating systems versions or distributions. In such cases, the containers may need to be updated or rebuilt.

To facilitate these tasks, we developed some tools to automatically start the OpenRAN Gym LXC containers on the different platforms considered in this work, and to properly interface them with the available radio resources. After the

<sup>9</sup>A detailed evaluation of OpenRAN Gym xApps, including their orchestration, and control of large-scale experimental networks can be found in [8, 21].

LXC images have been transferred (e.g., through the `scp` or `rsync` utilities) in a running instance of the testbed of interest, the image can be imported with the commands shown in Listing 6. This command imports the `scope-with-e2.tar.gz`

---

```

1 #!/bin/bash
2 lxc image import scope-with-e2.tar.gz --alias
   scope-e2
```

---

Listing 6: Commands to import the SCOPE LXC image with the E2 termination module.

LXC image transferred from Colosseum (i.e., the SCOPE image with the module for the E2 termination) to the compute machine of the remote testbed. After the above operation completes successfully, the new image, named `scope`, is visible by running the following command: `lxc image list`.

The LXC container can be, then, created from the imported image by running the commands shown in Listing 7.

---

```

1 #!/bin/bash
2 lxc init local:scope-e2 scope
```

---

Listing 7: Commands to create the SCOPE LXC container with the E2 termination module from the image imported in Listing 6.

After creating the container, additional operations may be required based on the specific OpenRAN Gym image, and SDR available in the remote testbed (e.g., USRP B210 or X310). As an example, if running the SCOPE container with an USRP B210, it is necessary to perform an *USB passthrough* operation to allow the container to use the USB interfaces and devices connected to the physical host (e.g., the USB interface to control the USRP). If using an USRP X310, instead, the container needs access to the network interface the host machine uses to communicate with the SDR, to set the right Maximum Transmission Unit (MTU) for it, and possibly to flash the FPGA of the USRP with the appropriate image. In both these cases, the container may require some additional permissions to be able to use the passed devices and interfaces (e.g., read/write permissions on the USB devices).

In the case of Colo-RAN—which can be imported and started with commands analogous to those shown in Listings 6 and 7—it is necessary to configure the Network Address Translation (NAT) rules of the host machine for it to forward of the messages directed to the RIC to the Colo-RAN container (e.g., the *E2 Setup Request* message used by the base station to subscribe to the RIC, and the *RIC Indication* messages used to send the KPMs to the xApps). Similarly to the previous case, the LXC container may require some additional permissions (e.g., to run the Docker containers of the Colo-RAN near-RT RIC in a *nested* manner inside the Colo-RAN LXC container).

After these operations have been executed, the LXC container (e.g., the SCOPE LXC container created in Listing 7) can be started with the commands of Listing 8. Now, the

---

```

1 #!/bin/bash
2 lxc start scope
```

---

Listing 8: Commands to start the SCOPE LXC container created in Listing 8.

OpenRAN Gym applications can be executed by following the procedures detailed in Sections III-A, IV-A, IV-B, and IV-C.

TABLE I: Compute node and radio setups used across the different testbeds.

Testbed	Compute Node	Processor	CPU Cores	RAM [GB]	Software-defined Radio
Base Station (BS) / UE					
Arena	Dell EMC PowerEdge R340	Intel Xeon E-2146G	6	32	NI USRP X310
Colosseum	Dell EMC PowerEdge R730	Intel Xeon E5-2650	48	128	NI USRP X310
COSMOS	Asus server (model undisclosed)	Intel i7-4790	4	16	NI USRP B210
POWDER (BS)	Dell EMC PowerEdge R740	Intel Xeon Gold 6126	24	98	NI USRP X310
POWDER (UE)	Intel NUC 8559	Intel 7-8559U	4	32	NI USRP B210
Near-RT RIC					
Arena	Dell EMC PowerEdge R340	Intel Xeon E-2146G	6	32	N/A
Colosseum	Dell EMC PowerEdge R730	Intel Xeon E5-2650	48	128	N/A
COSMOS	Supermicro 1028U-TRT+	Intel Xeon E5-2698	16	251	N/A
POWDER	Dell EMC PowerEdge R740	Intel Xeon Gold 6126	24	98	N/A

To simplify and automate the above setup operations, and to allow OpenRAN Gym users to swiftly configure and run the transferred containers, we developed and open-sourced a set of scripts that take care of (i) passing the right radio interface to the containers; (ii) giving them the required permissions; (iii) setting up the NAT rules of the host machine, and, finally, (iv) starting the OpenRAN Gym LXC containers from the imported images<sup>10</sup>. These scripts, which are supposed to be run after the commands of Listing 6, i.e., after the LXC image has been imported, are described in Listings 9 and 10.

Specifically, Listing 9 creates, sets up, and starts the SCOPE LXC container starting from the image imported in Listing 6. After creating the container on the testbed of interest (i.e.,

---

```
1 #!/bin/bash
2 ./start-lxc-scope.sh testbed usrp_type [flash]
```

---

Listing 9: Commands to start the SCOPE LXC container.

arena, powder, or cosmos), the script configures the USRP specified through the `usrp_type` parameter (i.e., b210 or x310) following the procedures described above (i.e., passing to the container the devices to interface with the USRP, and assigning the appropriate permissions to the container). The optional `flash` parameter also allows to flash the FPGA of the USRP X310 with the UHD image used by the container.<sup>11</sup>

The script of Listing 10 can be used to create, setup, and start the ColO-RAN near-RT RIC container starting from the image imported in Listing 6. This script creates the ColO-RAN

---

```
1 #!/bin/bash
2 ./start-lxc-ric.sh
```

---

Listing 10: Commands to start the ColO-RAN LXC container.

near-RT RIC container, assigns it the required permissions (e.g., to run the nested Docker containers), and starts it. Then, it sets the NAT rules of the host machine (where the container is running) for it to forward the messages intended for the RIC. Finally, it builds and starts the Docker containers of the ColO-RAN near-RT RIC (see Section IV) inside the created LXC container.

<sup>10</sup><https://github.com/wineslab/openrangegym-pawr>

<sup>11</sup>Please note that after the FPGA has been flashed with a new image, the USRP may need to be rebooted. We refer to the documentation of the various testbeds for the instructions on how to achieve this.

## VII. EXPERIMENTAL RESULTS

In this section, we showcase some experimental results obtained from running OpenRAN Gym and its components across a set of heterogeneous testbeds. We ported the SCOPE and ColO-RAN near-RT RIC containers from Colosseum to the Arena, POWDER, and COSMOS testbeds (see Sections II and VI). A description of the compute node and radio setups used in these testbeds (also summarized in Table I) follows. Since the capabilities offered by the different testbeds can be substantially different (e.g., number of available over-the-air nodes), for the sake of consistency, and to fairly compare results, we run experiments with one cellular base station and up to three UEs, and one near-RT RIC node. We divide the spectrum of the base stations into up to three network slices, and statically assign the UEs to them (e.g., based on the Service Level Agreement (SLA) between users and their network operator). Downlink User Datagram Protocol (UDP) traffic generated through the iPerf3 tool is leveraged to evaluate the network performance. Finally, the base stations—implemented through SCOPE—connect to ColO-RAN near-RT RIC through the E2 interface standardized by O-RAN.

**POWDER.** We instantiated both the ColO-RAN near-RT RIC and the SCOPE base station on Dell EMC PowerEdge R740 compute nodes with Intel Xeon Gold 6126 processor, 24 CPU cores and 98 GB memory. The UEs were instantiated on Intel NUC 8559 nodes with Intel i7-8559U processor, 4 CPU cores, and 32 GB RAM. The radio front-end of the base station was implemented through a USRP X310, while USRP B210 were used for the UEs. As this testbed does not natively support the LXC virtualization technology, the OpenRAN Gym container images were transferred from Colosseum to the compute nodes through the `scp` utility, instantiated on Ubuntu Linux images loaded on the bare-metal servers of the testbed.

**COSMOS.** In this case, the near-RT RIC was instantiated on a Supermicro 1028U-TRT+ server with an Intel Xeon E5-2698 processor, 16 CPU cores and 251 GB memory. Base station and UE, instead, were virtualized on Asus servers with Intel i7-4790 processor, 4 CPU cores, and 16 GB memory driving USRP B210 SDRs. Similarly to what done for POWDER, as the LXC virtualization technology is not directly supported by this testbed, the container images were transferred from Colosseum through the `scp` utility, and instantiated on Ubuntu Linux images loaded on the bare-metal nodes available on the

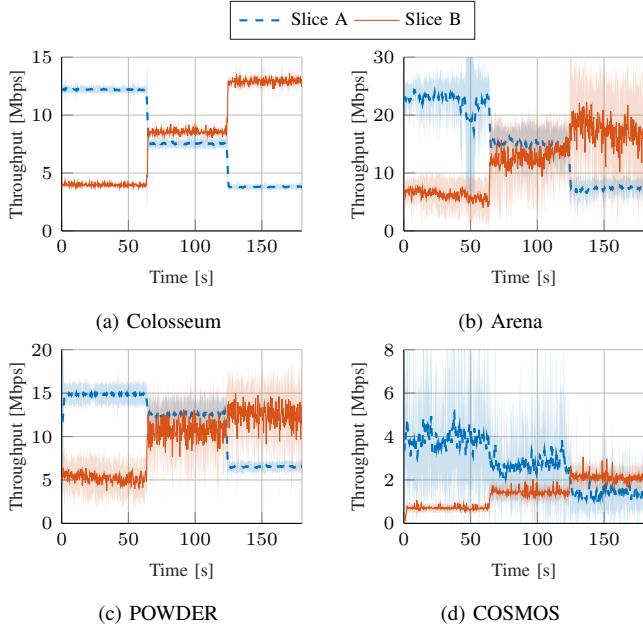


Figure 5: Overall slice throughput varying the percentage of RBGs allocated to each slice over time according to the configuration reported in Table II.

testbed.

**Arena.** All applications were run on Dell EMC PowerEdge R340 servers with Intel Xeon E-2146G processor, 6 CPU cores, and 32 GB memory. In this case, the OpenRAN Gym LXC containers are instantiated directly on the bare-metal nodes of the testbed, which leverage USRP X310 SDRs as radio front-ends. On this testbed, the UEs are implemented through commercial smartphones.

**Colosseum.** To mimic the same deployment scenario used in the other testbeds, in Colosseum we considered cellular nodes deployed in a static RF scenario without user mobility. In this case the LXC containers of RIC, base station, and UEs directly run on Colosseum bare-metal nodes, i.e., Dell EMC PowerEdge R730 servers with Intel Xeon E5-2650 processor, 48 CPU cores, and 128 GB memory. All the cellular nodes leverage USRP X310 SDRs as radio-front ends.

#### A. Results

To showcase the flexibility of OpenRAN Gym in dynamically reconfiguring the spectrum allocated to the network slices across different testbeds, Figures 5 and 6 show the overall throughput of each network slice varying the resources allocated to them, in terms of RBGs. 95% confidence intervals are also represented by the shaded areas in the figures. For both figures, the percentage of RBGs allocated to each slice of the base station—which uses a 10 MHz configuration—is dynamically changed through the SCOPE APIs according to the following configuration (also summarized in Table II). In Figure 5, the two network slices, i.e., slice A and B in the figure, are allocated the following RBGs percentage: (i) 75% to slice A and 25% to slice B in the first minute; (ii) 50% to each slice in the second minute, and (iii) 25% to slice A and 75% to slice B in the third minute. In Figure 6, instead they are allocated the following RBGs percentage: (i) 75% to slice A

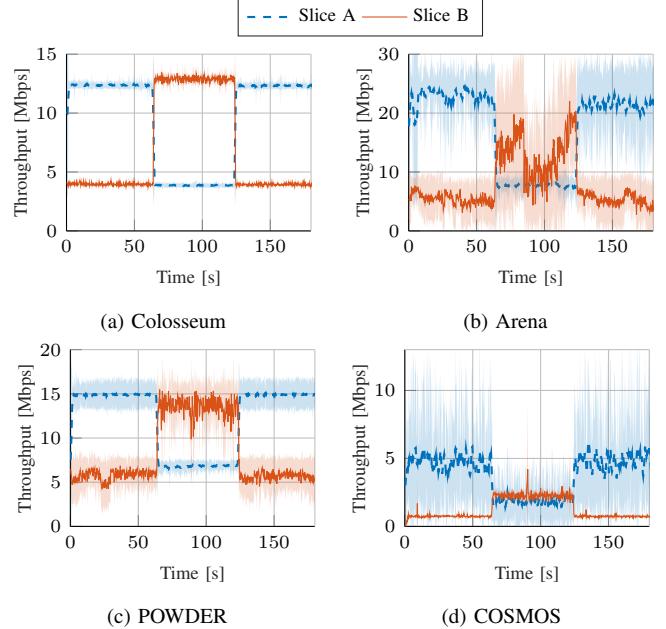


Figure 6: Overall slice throughput varying the percentage of RBGs allocated to each slice over time according to the configuration reported in Table II.

TABLE II: Slicing configuration, expressed as percentage of RBGs, used in Figures 5 and 6.

Figure	Slice	First Minute	Second Minute	Third Minute
Figure 5	Slice A	75% RBGs	50% RBGs	25% RBGs
	Slice B	25% RBGs	50% RBGs	75% RBGs
Figure 6	Slice A	75% RBGs	25% RBGs	75% RBGs
	Slice B	25% RBGs	75% RBGs	25% RBGs

and 25% to slice B in the first minute; (ii) 25% to slice A and 75% to slice B in the second minute, and (iii) 75% to slice A and 25% to slice B in the third minute. In both these figures, the throughput varies proportionally to the specific allocation of slice resources, in which slices with more RBGs achieve higher throughput values. These values then change during the experiment as RBGs are dynamically reallocated to the slices. We notice that even if the throughput differs across the various testbeds because of the different capabilities and environments they offer—with Arena achieving the highest performance due to the use of commercial smartphones as the UEs—the overall trends are consistent across the different setups.

We now showcase an instance in which the ColO-RAN near-RT RIC is leveraged to control a softwareized RAN implemented through SCOPE. LXC containers for both applications are deployed on the testbeds mentioned above, whose specifications are summarized in Table I. Figure 7 shows the evolution in time of the throughput of the three network slices (namely, slice A, B, and C) implemented by the SCOPE base station. Initially, the slices are allocated a fixed RBG configuration, and no control is performed by the RIC. Then, at around second 150, an xApp that prioritizes one of the network slices (slice A in the figure) is instantiated on the near-RT RIC. As a result, the xApp dynamically reallocates the amount of RBGs of each slice, which reflects on the

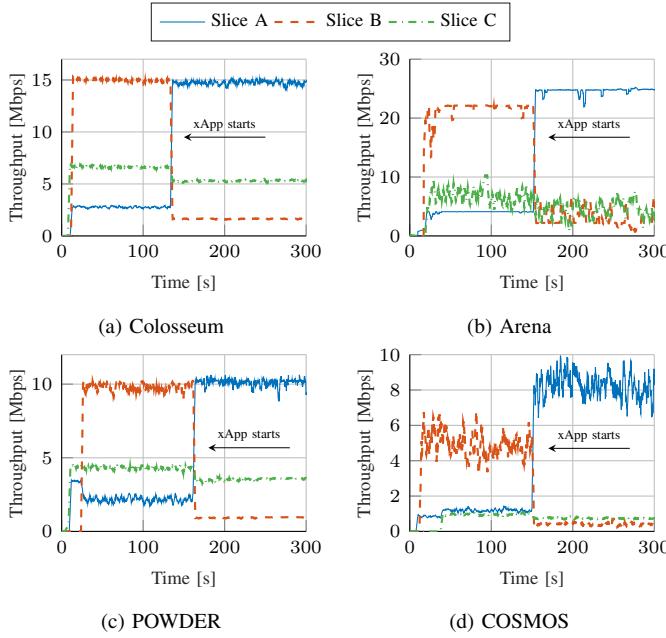


Figure 7: Slice throughput when the SCOPE RAN is controlled by ColO-RAN near-RT RIC. At around second 150, xApp to prioritize the amount of resources (i.e., RBGs) allocated to slice A is instantiated on the near-RT RIC.

performance of the slices of the RAN. Similar to the previous case, slices with a larger amount of RBGs allocated to them achieve higher throughput values. Overall, even in this case results are consistent across the different testbeds.

Now we show some timing statistics on the average amount of time taken to transfer the SCOPE and ColO-RAN LXC images from Colosseum to the Arena, COSMOS, and POWDER platforms. All the image transfers were performed through the `scp` utility, while the LXC containers were created following the procedures detailed in Section VI. In both cases, these timing statistics were derived using the hardware of Table I. We used the compute nodes listed in the “base station (BS)/UE” section of the table for the SCOPE LXC image/container (in the case of the POWDER platform, in which different compute nodes were listed for base station and UE, the base station node was used), and the compute nodes in “near-RT RIC” section of the table for ColO-RAN. In the tables that will be described next, we consider the following LXC images:

- *SCOPE w/ E2*: this is the SCOPE LXC image with the O-RAN E2 termination to interface with the near-RT RIC.
- *Colo-RAN near-RT RIC, prebuilt*: this is the ColO-RAN LXC image in which the Docker containers of the RIC, and sample xApp (described in Section IV) have been built a priori.
- *Colo-RAN near-RT RIC, to build*: this is the ColO-RAN LXC image with the scripts to build the Docker containers of the RIC and sample xApp from scratch.

Table III shows the average time required to transfer the LXC images from Colosseum to the other platforms. Times span from as low as  $\sim 1.5$  minutes to as high as almost 6 minutes, depending on the size of each image—also listed in the table—and capabilities of the testbeds. However, transfer times are consistent across the different testbeds.

TABLE III: Average time to transfer the LXC images from Colosseum to specific testbeds. The size of each image is listed in brackets.

Testbed	SCOPE w/ E2 (1.7 GB)	ColO-RAN near-RT RIC, prebuilt (6.5 GB)	ColO-RAN near-RT RIC, to build (1.6 GB)
Arena	1 m 27.413 s	5 m 41.487 s	1 m 25.002 s
COSMOS	1 m 28.631 s	5 m 39.704 s	1 m 27.352 s
POWDER	1 m 30.787 s	5 m 43.704 s	1 m 28.546 s

Finally, Table IV shows the times taken to instantiate LXC containers from the images transferred from Colosseum. In

TABLE IV: Average time to start as a container the LXC image exported from Colosseum on specific testbeds. The size of each image is listed in brackets.

Testbed	SCOPE w/ E2 (1.7 GB)	ColO-RAN near-RT RIC, prebuilt (6.5 GB)	ColO-RAN near-RT RIC, to build (1.6 GB)
Arena	0.887 s	1 m 11.483 s	46 m 18.110 s
COSMOS	25.463 s	2 m 34.905 s	26 m 4.410 s
POWDER	30.139 s	2 m 55.654 s	21 m 11.220 s

this case, we notice some difference among the times achieved on the different testbeds. For instance, Arena is significantly faster than COSMOS and POWDER in instantiating the SCOPE container—completing the instantiation in less than 1 s—and the prebuilt ColO-RAN container (instantiation in approximately 1 minute). This is mainly due to the fact that Arena allows users to instantiate applications on the bare-metal nodes directly. This removes the latency of the extra virtualization layer of the other two testbeds, in which the LXC containers are nested inside the virtualized architecture the users are given access to. When it comes to building the Docker containers of the ColO-RAN near-RT RIC (see Section IV) from scratch, instead, POWDER and COSMOS are significantly faster than Arena, taking approximately half the time to complete the same operations. This is mainly due to the superior compute capabilities of the nodes of these two testbeds (24 core CPU server on POWDER, and 16 core server on COSMOS vs. 6 core CPU server on Arena). Nonetheless, this building operation needs to be completed only once, as the compiled ColO-RAN LXC image can be saved to be used in subsequent experiments, with instantiation times sensibly lower (slightly above 1 minute for Arena, and below 3 minutes for POWDER and COSMOS).

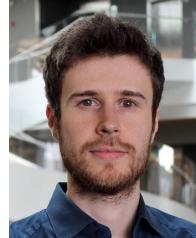
## VIII. CONCLUSIONS

We presented OpenRAN Gym, the first publicly-available research platform for data-driven O-RAN experimentation at scale on heterogeneous wireless testbeds. Building on, and extending, frameworks for data collection and RAN control, OpenRAN Gym enables the end-to-end design and testing of data-driven xApps instantiated on the O-RAN infrastructure. We described the core components of OpenRAN Gym—including frameworks and experimental platforms—and detailed procedures and configuration options for experimenting at scale on a softwarized RAN instantiated on Colosseum. Then, we gave an overview of the xApp design and testing workflow enabled by OpenRAN Gym, also showcasing an example of two xApps used to control a large-scale O-RAN managed softwarized RAN deployed on Colosseum. Finally, we

demonstrated how OpenRAN Gym solutions and experiments can be transitioned from Colosseum to heterogeneous real-world platforms, such as the Arena testbed, and the POWDER and COSMOS platforms of the PAWR program. OpenRAN Gym is publicly-available to the research community, and opened up for community contributions and additions.

## REFERENCES

- [1] L. Bonati, M. Polese, S. D’Oro, S. Basagni, and T. Melodia, “OpenRAN Gym: An Open Toolbox for Data Collection and Experimentation with AI in O-RAN,” in *Proceedings of IEEE WCNC Workshop on Open RAN Architecture for 5G Evolution and 6G*, Austin, TX, USA, April 2022.
- [2] 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, pp. 1–28, December 2020.
- [3] O-RAN Working Group 1, “O-RAN Architecture Description 5.00,” O-RAN.WG1.O-RAN-Architecture-Description-v05.00 Technical Specification, July 2021.
- [4] M. Polese, L. Bonati, S. D’Oro, S. Basagni, and T. Melodia, “Understanding O-RAN: Architecture, Interfaces, Algorithms, Security, and Research Challenges,” *arXiv:2202.01032 [cs.NI]*, February 2022.
- [5] O-RAN Working Group 3, “O-RAN Near-RT RAN Intelligent Controller Near-RT RIC Architecture 2.00,” O-RAN.WG3.RICARCH-v02.00, March 2021.
- [6] O-RAN Working Group 2, “O-RAN Non-RT RIC Architecture 1.0,” O-RAN.WG2.Non-RT-RIC-ARCH-TS-v01.00 Technical Specification, July 2021.
- [7] ———, “O-RAN AI/ML workflow description and requirements 1.03,” O-RAN.WG2.AIML-v01.03 Technical Specification, July 2021.
- [8] M. Polese, L. Bonati, S. D’Oro, S. Basagni, and T. Melodia, “ColO-RAN: Developing Machine Learning-based xApps for Open RAN Closed-loop Control on Programmable Experimental Platforms,” *IEEE Transactions on Mobile Computing*, pp. 1–14, July 2022.
- [9] S. D’Oro, M. Polese, L. Bonati, H. Cheng, and T. Melodia, “dApps: Distributed Applications for Real-time Inference and Control in O-RAN,” *IEEE Communications Magazine*, pp. 1–7, 2022, in print; preprint available at <https://arxiv.org/pdf/2203.02370.pdf>.
- [10] L. Bonati, P. Johari, M. Polese, S. D’Oro, S. Mohanti, M. Tehrani-Moayyed, D. Villa, S. Shrivastava, C. Tassie, K. Yoder, A. Bagga, P. Patel, V. Petkov, M. Seltser, F. Restuccia, A. Gosain, K. R. Chowdhury, S. Basagni, and T. Melodia, “Colosseum: Large-Scale Wireless Experimentation Through Hardware-in-the-Loop Network Emulation,” in *Proceedings of IEEE DySPAN*, December 2021.
- [11] L. Bonati, S. D’Oro, S. Basagni, and T. Melodia, “SCOPE: An open and softwareized prototyping platform for NextG systems,” in *Proceedings of ACM MobiSys*, June 2021.
- [12] L. Bertizzolo, L. Bonati, E. Demirors, A. Al-Shawabka, S. D’Oro, F. Restuccia, and T. Melodia, “Arena: A 64-antenna SDR-based Ceiling Grid Testing Platform for Sub-6 GHz 5G-and-Beyond Radio Spectrum Research,” *Computer Networks*, vol. 181, pp. 1–17, November 2020.
- [13] Platforms for Advanced Wireless Research (PAWR). <https://www.advancedwireless.org>. Accessed December 2021.
- [14] J. Breen, A. Buffinire, J. Duerig, K. Dutt, E. Eide, A. Ghosh, M. Hibler, D. Johnson, S. K. Kasera, E. Lewis *et al.*, “POWDER: Platform for Open Wireless Data-driven Experimental Research,” *Computer Networks*, vol. 197, pp. 1–18, October 2021.
- [15] D. Raychaudhuri, I. Seskar, G. Zussman, T. Korakis, D. Kilper, T. Chen, J. Kolodziejski, M. Sherman, Z. Kostic, X. Gu, H. Krishnaswamy, S. Moshwari, P. Skrimponis, and C. Guterman, “Challenge: COSMOS: A City-Scale Programmable Testbed for Experimentation with Advanced Wireless,” in *Proceedings of ACM MobiCom*, London, United Kingdom, Sept. 2020.
- [16] M. Dryjanski, L. Kulacz, and A. Kliks, “Toward Modular and Flexible Open RAN Implementations in 6G Networks: Traffic Steering Use Case and O-RAN xApps,” *Sensors*, vol. 21, no. 24, pp. 1–14, December 2021.
- [17] D. Johnson, D. Maas, and J. Van Der Merwe, “Open Source RAN Slicing on POWDER: A Top-to-Bottom O-RAN Use Case,” in *Proceedings of ACM MobiSys*, June 2021.
- [18] H. Lee, J. Cha, D. Kwon, M. Jeong, and I. Park, “Hosting AI/ML Workflows on O-RAN RIC Platform,” in *Proceedings of IEEE GLOBECOM Workshops*, December 2020.
- [19] A. S. Abdalla, P. S. Upadhyaya, V. K. Shah, and V. Marojevic, “Toward Next Generation Open Radio Access Network—What O-RAN Can and Cannot Do!” *arXiv preprint arXiv:2111.13754 [cs.NI]*, November 2021.
- [20] O-RAN Alliance Conducts First Global Plugfest to Foster Adoption of Open and Interoperable 5G Radio Access Networks. (2019, December) <https://tinyurl.com/f48auynf>.
- [21] S. D’Oro, L. Bonati, M. Polese, and T. Melodia, “OrchestraRAN: Network Automation through Orchestrated Intelligence in the Open RAN,” in *Proceedings of IEEE INFOCOM*, May 2022, arXiv:2201.05632 [cs.NI].
- [22] I. Gomez-Miguelez, A. Garcia-Saavedra, P. Sutton, P. Serrano, C. Cano, and D. Leith, “srsLTE: An open-source platform for LTE evolution and experimentation,” in *Proceedings of ACM WiTECH*, October 2016.
- [23] F. Kaltenberger, A. P. Silva, A. Gosain, L. Wang, and T.-T. Nguyen, “OpenAirInterface: Democratizing innovation in the 5G era,” *Computer Networks*, no. 107284, May 2020.
- [24] M. Kohli, T. Chen, M. B. Dastjerdi, J. Welles, I. Seskar, H. Krishnaswamy, and G. Zussman, “Open-Access Full-Duplex Wireless in the ORBIT and COSMOS Testbeds,” *Computer Networks*, 2021.
- [25] L. Bonati, S. D’Oro, M. Polese, S. Basagni, and T. Melodia, “Intelligence and Learning in O-RAN for Data-driven NextG Cellular Networks,” *IEEE Communications Magazine*, vol. 59, no. 10, pp. 21–27, October 2021.
- [26] O-RAN Software Community. O-DU L2 Repository. <https://github.com/o-ran-sc/o-du-l2>. Accessed December 2021.
- [27] David Johnson. POWDER RIC Profile Repository. <https://gitlab.flux.utah.edu/johnson/ric-profile>. Accessed December 2021.



**Leonardo Bonati** received his B.S. in Information Engineering and his M.S. in Telecommunication Engineering from University of Padova, Italy in 2014 and 2016, respectively. He is currently pursuing a Ph.D. degree in Computer Engineering at Northeastern University, MA, USA. His research interests focus on 5G and beyond cellular networks, network slicing, and software-defined networking for wireless networks.



**Michele Polese** is a Principal Research Scientist at the Institute for the Wireless Internet of Things, Northeastern University, Boston, since March 2020. He received his Ph.D. at the Department of Information Engineering of the University of Padova in 2020. He also was an adjunct professor and postdoctoral researcher in 2019/2020 at the University of Padova, and a part-time lecturer in Fall 2020 and 2021 at Northeastern University. During his Ph.D., he visited New York University (NYU), AT&T Labs in Bedminster, NJ, and Northeastern University. His

research interests are in the analysis and development of protocols and architectures for future generations of cellular networks (5G and beyond), in particular for millimeter-wave and terahertz networks, spectrum sharing and passive/active user coexistence, open RAN development, and the performance evaluation of end-to-end, complex networks. He has contributed to O-RAN technical specifications and submitted responses to multiple FCC and NTIA notice of inquiry and requests for comments, and is a member of the Committee on Radio Frequency Allocations of the American Meteorological Society (2022-2024). He collaborates and has collaborated with several academic and industrial research partners, including AT&T, Mavenir, NVIDIA, InterDigital, NYU, University of Aalborg, King’s College, and NIST. He was awarded with several best paper awards, is serving as TPC co-chair for WNS3 2021-2022, as an Associate Technical Editor for the IEEE Communications Magazine, and has organized the Open 5G Forum in Fall 2021. He is a Member of the IEEE.



of the IEEE.

**Salvatore D'Oro** is a Research Assistant Professor at Northeastern University. He received his Ph.D. degree from the University of Catania in 2015. Salvatore is an area editor of Elsevier Computer Communications journal and serves on the Technical Program Committee (TPC) of multiple conferences and workshops such as IEEE INFOCOM, IEEE CCNC, IEEE ICC and IFIP Networking. Dr. D'Oro's research interests include optimization, artificial intelligence, security, network slicing and their applications to 5G networks and beyond. He is a Member



**Stefano Basagni** is with the Institute for the Wireless Internet of Things and a professor at the ECE Department at Northeastern University, in Boston, MA. He holds a Ph.D. in electrical engineering from the University of Texas at Dallas (2001) and a Ph.D. in computer science from the University of Milano, Italy (1998). Dr. Basagni's current interests concern research and implementation aspects of mobile networks and wireless communications systems, wireless sensor networking for IoT (underwater, aerial and terrestrial), and definition and performance evaluation of network protocols. Dr. Basagni has published over ten dozen of highly cited, refereed technical papers and book chapters. His h-index is currently 46 (May 2022). He is also co-editor of three books. Dr. Basagni served as a guest editor of multiple international ACM/IEEE, Wiley and Elsevier journals. He has been the TPC co-chair of international conferences. He is a distinguished scientist of the ACM, a senior member of the IEEE, and a member of CUR (Council for Undergraduate Education).



**Tommaso Melodia** is the William Lincoln Smith Chair Professor with the Department of Electrical and Computer Engineering at Northeastern University in Boston. He is also the Founding Director of the Institute for the Wireless Internet of Things and the Director of Research for the PAWR Project Office. He received his Ph.D. in Electrical and Computer Engineering from the Georgia Institute of Technology in 2007. He is a recipient of the National Science Foundation CAREER award. Prof. Melodia has served as Associate Editor of IEEE Transactions on Wireless Communications, IEEE Transactions on Mobile Computing, Elsevier Computer Networks, among others. He has served as Technical Program Committee Chair for IEEE Infocom 2018, General Chair for IEEE SECON 2019, ACM Nanocom 2019, and ACM WUWNet 2014. Prof. Melodia is the Director of Research for the Platforms for Advanced Wireless Research (PAWR) Project Office, a \$100M public-private partnership to establish 4 city-scale platforms for wireless research to advance the US wireless ecosystem in years to come. Prof. Melodia's research on modeling, optimization, and experimental evaluation of Internet-of-Things and wireless networked systems has been funded by the National Science Foundation, the Air Force Research Laboratory the Office of Naval Research, DARPA, and the Army Research Laboratory. Prof. Melodia is a Fellow of the IEEE and a Senior Member of the ACM.