

# Colosseum: Large-Scale Wireless Experimentation Through Hardware-in-the-Loop Network Emulation

Leonardo Bonati,<sup>\*</sup> Pedram Johari,<sup>\*</sup> Michele Polese,<sup>\*</sup> Salvatore D’Oro,<sup>\*</sup> Subhramoy Mohanti,<sup>\*</sup> Miead Tehrani-Moayyed,<sup>\*</sup> Davide Villa,<sup>\*</sup> Shweta Shrivastava,<sup>\*</sup> Chinenye Tassie,<sup>\*</sup> Kurt Yoder,<sup>†</sup> Ajeet Bagga,<sup>‡</sup> Paresh Patel,<sup>‡</sup> Ventz Petkov,<sup>‡</sup> Michael Seltser,<sup>‡</sup> Francesco Restuccia,<sup>\*§</sup> Abhimanyu Gosain,<sup>\*</sup> Kaushik R. Chowdhury,<sup>\*</sup> Stefano Basagni,<sup>\*</sup> Tommaso Melodia<sup>\*</sup>

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

<sup>§</sup>Roux Institute, Northeastern University, Portland, ME, U.S.A.

<sup>†</sup>Greenfly SAU LLC, Dunn Loring, VA, U.S.A. <sup>‡</sup>Cerbo IO LLC, Burlington, MA, U.S.A.

**Abstract**—Colosseum is an open-access and publicly-available large-scale wireless testbed for experimental research via virtualized and softwarized waveforms and protocol stacks on a fully programmable, “white-box” platform. Through 256 state-of-the-art software-defined radios and a massive channel emulator core, Colosseum can model virtually any scenario, enabling the design, development and testing of solutions at scale in a variety of deployments and channel conditions. These Colosseum radio-frequency scenarios are reproduced through high-fidelity FPGA-based emulation with finite-impulse response filters. Filters model the taps of desired wireless channels and apply them to the signals generated by the radio nodes, faithfully mimicking the conditions of real-world wireless environments. In this paper, we introduce Colosseum as a testbed that is for the first time open to the research community. We describe the architecture of Colosseum and its experimentation and emulation capabilities. We then demonstrate the effectiveness of Colosseum for experimental research at scale through exemplary use cases including prevailing wireless technologies (e.g., cellular and Wi-Fi) in spectrum sharing and unmanned aerial vehicle scenarios. A roadmap for Colosseum future updates concludes the paper.

**Index Terms**—Experimental Wireless Research, Wireless Channel Emulation, Data Factory, Mobile Networks, Artificial Intelligence.

## I. INTRODUCTION

Wireless technologies are undergoing advances at a pace never seen before. The 5th generation (5G) of cellular networks is evolving toward open and disaggregated deployments in a way never experienced by previous generations [1]; Wi-Fi standards are multiplying and now occupy new portions of the electromagnetic spectrum [2]; applications of wireless networking include new verticals, ranging from those involving Unmanned Aerial Vehicles (UAVs) [3, 4] to telemedicine [5].

The wireless research community has increasingly recognized the need for programmable, large-scale experimental testbeds to validate new enabling technologies and application scenarios in realistic environments, at scale, and with a diversity of traffic, topology, and channel conditions. Notoriously, however, high-fidelity testbeds are not easily accessible, scale is often limited, and costs are high. As a result, innovative solutions in wireless are often tested either through numerical network simulations (that often lack fidelity, especially as concerns modeling of hardware constraints and imperfections, or of the

effects of interference and channel variations on applications), or in small laboratory setups (which are able to capture to some extent the characteristics of hardware devices of commercial deployments but not the scale of such deployments or real propagation environments and interference conditions).

To bridge this gap, a number of high-profile programs have attempted to develop shared, community-driven, and remotely-accessible experimental wireless platforms. A common goal of these initiatives is to create a widely-accepted and validated ecosystem for repeatable and reproducible experimentation in the wireless networking field. For example, in the United States, the National Science Foundation (NSF) has spearheaded the Platforms for Advanced Wireless Research (PAWR) program. The goal of PAWR is to develop four city-scale programmable wireless platforms open to the research community at large. Among these, POWDER-RENEW focuses on developing a sub-6 GHz programmable cellular architecture [6] with massive MIMO capabilities [7]; COSMOS focuses on a combination of mmWave links [8] with low-latency programmable optical backhauls; AERPAW focuses on the intersection between UAV verticals and cellular applications [9]; while ARA will be focused on wireless rural scenarios [10].

Similar efforts are being carried out in Europe by the 5GENESIS project [11], which encompasses testbeds comprising an edge-computing Network Function Virtualization (NFV)-enabled radio infrastructure, orchestration and management frameworks, satellite communications, and ultra-dense network deployments. Although the testbeds of this program complement each other in offering the tools to evaluate several aspects of the 5G ecosystem, they ignore non-cellular wireless access technologies. Smaller-scale indoor testbeds include the Drexel Grid [12], ORBIT [8] and Arena [13]. The Drexel Grid couples physical Software-defined Radios (SDRs) with virtual nodes capable of emulating custom channel conditions. ORBIT comprises a grid of SDRs and compute nodes equipped with transceivers for heterogeneous wireless technologies (e.g., Bluetooth, ZigBee and LTE). Finally, Arena mounts a ceiling grid of 64 SDR-driven antennas connected with edge compute and Artificial Intelligence (AI) capabilities. While these testbeds provide a good representation of indoor wireless propagation environments, their scale can hardly capture the dynamics of large real-world deployments.

In this paper, we present Colosseum—the world’s largest

Contact: melodia@northeastern.edu. This work was partially supported by the U.S. National Science Foundation under grant CNS-1925601.




Fig. 1: Colosseum architecture.

wireless network emulator with hardware in-the-loop—as a platform that is *for the first time* available to the research community.<sup>1</sup> Originally built by the Defense Advanced Research Projects Agency (DARPA) and by the Johns Hopkins University Applied Physics Laboratory to support the Spectrum Collaboration Challenge (SC2) [14–24], Colosseum is being expanded and operated by the Institute for the Wireless Internet of Things<sup>2</sup> at Northeastern University through an NSF grant, which also made it publicly available to the research community. With its 256 SDRs and 128 remotely-accessible compute nodes and GPUs, Colosseum provides the capabilities to test full-protocol stack solutions at scale with real hardware devices and in emulated—yet realistic—environments with complex channel interactions (e.g., path loss, fading, multipath). Besides its experimentation capabilities, Colosseum can be used as an AI playground and wireless data factory to create large-scale datasets and train/test solutions in a safe and controlled environment [25, 26]. We provide an overview of the architectural components and emulation capabilities of Colosseum in Sections II and III. We show how to use Colosseum in Section IV, including practical use cases in Section V. Finally, we describe the planned extensions to Colosseum in Section VI and we draw our conclusions in Section VII.

## II. COLOSSEUM ARCHITECTURE

A high-level representation of the architecture of Colosseum is shown in Fig. 1. Colosseum comprises 128 Standard Radio Nodes (SRNs), the Massive Channel Emulator (MCHEM), the Radio Frequency (RF) scenario server, the Traffic Generator (TGEN), and the management infrastructure. The SRNs, which can be controlled remotely to perform experiments, are divided in four quadrants and are synchronized in time and frequency through hierarchical OctoClock clock distributors. Each SRN is a state-of-the-art server with 48-core Intel Xeon E5-2650 CPUs and an NVIDIA Tesla K40m GPU, and drives an NI/Ettus USRP X310. Each X310 is equipped with two UBX-160 daughterboards that operate between 10 MHz and 6GHz. Colosseum allows multiple concurrent users to automatically deploy software containers—implemented via Linux Containers (LXCs)—on the bare-metal SRNs (Fig. 1, left).

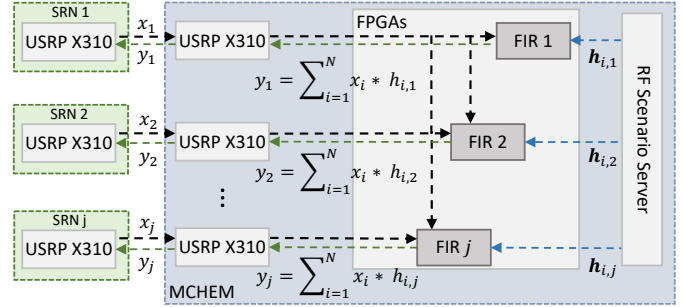


Fig. 2: FPGA-based RF scenario emulation in Colosseum.

Containers can be used to run a variety of protocol stacks and to control parameters and configurations at different layers of these stacks.

MCHEM is tasked with the channel emulation in Colosseum. It comprises four interconnected quadrants, each with 4 NI ATCA 3671 FPGA modules and 16 Virtex-7 690T FPGAs, and it drives an array of 128 USRPs X310 interconnected in a one-to-one fashion to the USRPs of the SRNs (see Fig. 2).

When an RF transmission occurs in Colosseum, the signals generated by the USRPs at the SRN side (e.g., signal  $x_1$  in Fig. 2) get transmitted to the corresponding USRPs X310 of MCHEM, which perform RF to baseband and analog-to-digital conversions. The digital signals are then forwarded to the FPGAs of MCHEM that process them through Finite Impulse Response (FIR) filters. Filters are composed of 512 pre-computed complex-valued taps that capture the characteristics of the channel, i.e., the Channel Impulse Response (CIR), between any pair of SRNs.<sup>3</sup> As depicted in Fig. 2, the FIR filters load the vector  $\mathbf{h}_{i,j}$  of the 512-tap CIRs among receive node  $j$  and every transmit node  $i$ , with  $i, j \in \{1, \dots, N\}$  set of SRNs of a specific experiment. These channel taps are then applied to signal  $x_i$  through a convolution operation. The resulting signal  $y_j = \sum_{i=1}^N x_i * h_{i,j}$  (i.e., the transmitted signals  $x_i$  processed with the CIR of the emulated channel between nodes  $i$  and  $j$ ) is then sent to SRN  $j$ . In this way, Colosseum aggregates all the received signals after passing them through the corresponding CIR and transmits the results  $y_j$  to all the SRNs of the experiment (and not only to the intended receiver),

<sup>1</sup><https://www.colosseum.net>

<sup>2</sup><https://www.northeastern.edu/wiot>

<sup>3</sup>Due to the high computational complexity required to generate scenarios, and to the large space to store them, only 4 channel taps are non-zero-valued.

thus capturing and emulating effects of real wireless channels, such as the interference among nodes and superimposition of signals from multiple nodes [27].

The RF scenario server (see Fig. 1 and 2) maintains a catalog of the Colosseum RF scenarios and feeds their channel taps to the channel emulator at run time. Scenarios make it possible to emulate effects of the wireless channel, including path loss and fading in terrains up to 1km<sup>2</sup> and with up to 80MHz bandwidth. The modular architecture of MCHM and the independence of its USRPs allow Colosseum to concurrently emulate different scenarios on different experiments so that multiple users can operate on the system at the same time. As we will discuss in Section IV, the specific scenario to run can be selected by the user through a specialized control interface.

Similarly to how MCHM emulates the RF scenarios, TGEN—based on the U.S. Naval Research Laboratory’s Multi-Generator (MGEN) [28]—takes care of emulating IP traffic flows between the SRNs. The set of traffic flows that can be generated by TGEN, together with their characteristics (e.g., packet rate, size and distribution), constitutes a *traffic scenario*. Once a traffic scenario starts, packets are delivered to the SRNs, which handle them through the user-defined protocol stack (e.g., by transmitting them through a cellular or Wi-Fi stack, see Section V). Finally, the Colosseum management infrastructure hosts a variety of auxiliary services (Fig. 1). These include: (i) the website through which the users reserve resources and start experiments; (ii) the resource manager, which allocates resources to the users; (iii) gateways for user and management access to Colosseum; (iv) a 900 TB Network Attached Storage (NAS) to store LXC images, experiment data and logs, and (v) various network services, including services to provide time synchronization across the whole testbed.

### III. EXPERIMENTAL SCENARIOS IN COLOSSEUM

Experimental scenarios, i.e., RF and traffic scenarios, are at the core of Colosseum. RF scenarios enable the emulation of up to 65,535 wireless channels in diverse environments with an emulation area up to 1 km<sup>2</sup>. They give users full control over the wireless environment, allowing them to capture and reproduce the desired channel effects while enabling reproducibility and repeatability of experiments at scale. This emulation is carried out by MCHM through FPGA-based FIR filters, which capture and reproduce the CIR of the emulated wireless environment (see Section II). The CIR can be derived in many different ways: for instance, it can be computed through well-known channel modeling equations, performing field measurements [29], or leveraging high-accuracy software tools, such as ray-tracers [30]. In this way, scenarios not only model the channel between transmitter and receiver pairs, but they also model effects typical of the wireless propagation environment, such as interference and superposition of the signals generated by multiple nodes. This ultimately guarantees a *high-fidelity* channel emulation, ensuring that the emulated environments are as close as possible to real-world deployments.

#### A. Examples of Available Experimental Scenarios

We now provide a sample of large-scale scenarios that are available on Colosseum at the time of this writing: the *Alleys of Austin* and the *SC2 Championship Event (SCE) Qualification* scenarios—developed by DARPA for SC2—and a set of *cellular scenarios* developed as part of [25].

**Alleys of Austin.** This scenario is set in the urban area of downtown Austin, TX, and involves 50 nodes divided in 5 groups, or *squads*. Every group consists of 9 pedestrian users moving in a row formation, and a UAV circling on top of them. The center frequency of this scenario is set to 1 GHz. The duration of the scenario is 15 minutes, divided in three stages of 5 minutes each. In the first stage, nodes exchange voice traffic. In the second stage, they transmit images and videos in addition to the voice data of the previous phase. Finally, in the third stage, the transmission rate of the nodes significantly increases.

**SCE Qualification.** This scenario concerns 10 nodes and has center frequency set to 1 GHz. The duration of the scenario is 10 minutes. The Signal-to-Noise-Ratio (SNR) among nodes is initially set to 20 dB and decreases by 5 dB every 2 minutes of the scenario. Finally, in the last 2 minutes, the scenario center frequency shifts to 1.1 GHz, and the SNR is reset to 20 dB. As for the traffic, nodes exchange UDP packets at constant bitrate.

**Cellular Scenarios.** Colosseum includes a number of scenarios designed for cellular networking experimentation developed as part of [25]. They involve 8 to 10 base stations serving four mobile users each. The locations of the base stations—derived from the OpenCellID database [31]—match those of real-world cellular deployments in Rome, Italy, in Boston, MA, and in the POWDER testbed [6] of Salt Lake City, UT. The central frequency of these scenarios is set to 1 GHz and they have a duration of 10 minutes each. Users are deployed at different distances from the base stations, i.e., closer or farther from them, and they either move with pedestrian speed or remain still for the whole scenario. Finally, the uplink and downlink traffic profiles among users and base stations entail video packets for video streaming applications.

#### B. Colosseum as a Wireless Data Factory

Colosseum scale and emulation capabilities allow it to create large-scale datasets, namely acting as a wireless data factory. Its unique emulation capabilities—both in terms of RF and traffic scenarios—allow users to configure an experiment once and then seamlessly run it in different environment, channel and traffic conditions. Moreover, once set up, experiments can be run automatically in *batches*, allowing users to collect data in bulk into wireless datasets (see Section IV). These capabilities are fundamental when designing Machine Learning (ML) algorithms, which need to be trained on large amounts of data and on different conditions to be able to generalize and adapt to unforeseen situations.

### IV. OPERATIONAL MODES OF COLOSSEUM

Colosseum allows users to operate in either (i) *interactive mode*, in which users operate the SRNs manually, or (ii) *batch mode*, in which experiments are carried out automatically.

**Interactive Mode.** The workflow of an interactive experiment in Colosseum is shown in Fig. 3. As a first step, the users

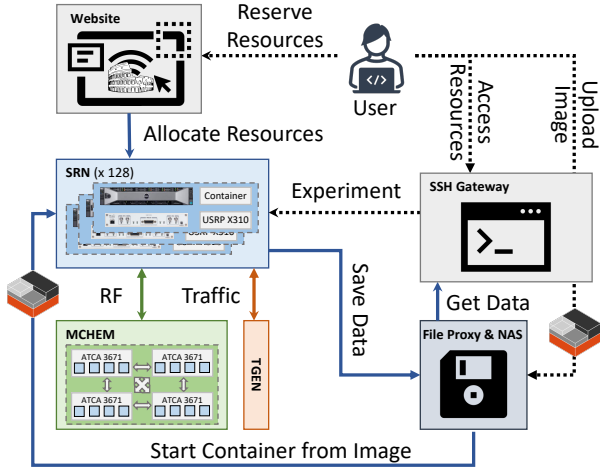


Fig. 3: Workflow of a Colosseum interactive experiment.

reserve a number of SRNs through Colosseum website. In this step, they can specify which LXC image to load on each SRN (i.e., either one of Colosseum-provided images or an image they previously created/uploaded to the system). To guarantee a fair use of the testbed, Colosseum implements a token-based resource allocation system, where occupying resources costs a certain amount of tokens per hour.<sup>4</sup>

After resources have been reserved, Colosseum automatically instantiates LXC containers on the SRNs. Next, users log into Colosseum gateway through the Secure Shell (SSH) protocol and access the SRNs allocated to them. This is where the main part of Colosseum interactive experiments takes place. Users can start RF and traffic scenarios by using `colosseumcli`—an Application Programming Interface (API) to interface with some auxiliary services of Colosseum, e.g., MCHEM and TGEN—and control the USRPs X310 through software-defined protocol stacks instantiated on the SRNs (see Section V for exemplary Colosseum use cases). Through `colosseumcli`, users can also make a *snapshot* of the container running on the SRNs. This creates an LXC image from the container and saves it on Colosseum NAS for later use. Finally, when the experiment ends, resources are deallocated and the content of the `/logs` directory of each container are saved on the NAS, from where they can be retrieved through the file proxy server.

**Batch Mode.** While the interactive mode is more suitable for designing, prototyping and troubleshooting solutions, the batch mode can be used to automatically run experiments in bulk, e.g., to benchmark solutions or perform data collection. Batch jobs are configured through `json`-formatted files, in which users can specify the experiment duration, which RF and traffic scenarios to run, the number of SRNs to allocate and which LXC images to instantiate on them. Additional parameters—

<sup>4</sup>Colosseum users are divided into teams, i.e., sets of users belonging to the same research group or organization. Each team is provided with a token budget that automatically resets on a weekly basis.

handled by the user-defined programs—can also be passed. Once scheduled, batch jobs are inserted in a queue and kicked off based on resource availability. Once a batch job starts, Colosseum takes care of instantiating the containers and to run the specified emulation scenario. Then, user-defined startup scripts are invoked on the SRNs and the actual experiment is carried out. Once the batch job ends, data and logs are copied from the `/logs` directory of the SRN containers to the NAS, analogously to what happens in interactive mode.

## V. USE CASES OF COLOSSEUM

This section illustrates a set of Colosseum use cases. Examples concerning cellular networking are shown in Section V-A, applications using a software-defined Wi-Fi protocol stack are showcased in Section V-B, while spectrum sharing capabilities are discussed in Section V-C. Finally, the application of UAVs to wireless networking is demonstrated in Section V-D.

### A. Cellular Networking

Softwarization and virtualization will play a fundamental role in 5G and beyond cellular networks. Telecom operators will deploy open source custom solutions on a “white-box” agnostic Radio Access Network (RAN) where services, e.g., the base stations, will be instantiated on-demand [1, 32]. This not only endows the network with flexibility by design, but it also allows real-time optimization of the users’ service based on the current network conditions and traffic demand [26, 33]. Although open-source approaches bring unprecedented performance improvements, they require a reliable development environment and at scale evaluation before they are deployed on the commercial infrastructure. More importantly, this testing phase needs to account for different environments, wireless channel and traffic conditions. With its unique emulation capabilities, Colosseum allows the research community to reliably prototype and benchmark solutions for future cellular networks. To showcase these capabilities, we instantiate a cellular network with 4 base stations and 24 users through `srsRAN` [34]. Figure 4 shows the network downlink throughput when the users request traffic through `iperf3` at different rates: low traffic (1 Mbps per user), medium traffic (2 Mbps), and high traffic (4 Mbps). In this case, the base stations use a 10 MHz channel bandwidth.




Fig. 4: Downlink throughput with different traffic conditions.

Figure 5 shows how Colosseum scenarios can be used to study the impact of deployment distance and user mobility in cellular networks [25]. We evaluate a network with 10 base stations and 40 users in close proximity to the base stations. Users request downlink video traffic at a rate of 1 Mbps through




Fig. 5: Downlink spectral efficiency of the network varying the distance between users and base stations, and speed of the users.

one of Colosseum traffic scenarios. Specifically, Fig. 5a shows the downlink spectral efficiency of the network for different deployment distances between users and base stations: close (users are deployed within 20 m from the base stations), medium (50 m), and far (100 m). Figure 5b, instead, depicts the same metric for different configurations of user mobility: static (no mobility), moderate (users move at a speed of 3 m/s), and fast (5 m/s). In this case, the users are deployed in close proximity to the base stations. We notice that although the spectral efficiency decreases with the increased mobility/distance between users and base stations, the network performance shows tight 95% confidence intervals (shaded areas). This confirms the reliability and repeatability of Colosseum emulation.

**O-RAN.** Another use case of interest to cellular networks is that of the Open RAN, which envisions the deployment of open source and softwarized stacks and components on a “white-box” hardware infrastructure. In this regard, the standardization efforts of the O-RAN Alliance are worth mentioning. This consortium is working at the standardization of open interfaces among the white-box network components—which would enable interoperability of equipment from multiple vendors—and at the creation of open controllers for the cellular infrastructure, namely RAN Intelligent Controllers (RICs). These controllers operate at different timescales to govern the components of the Open RAN and enable the execution of third-party applications, e.g., xApps, that control the RAN elements (see Fig. 6) [1, 26].

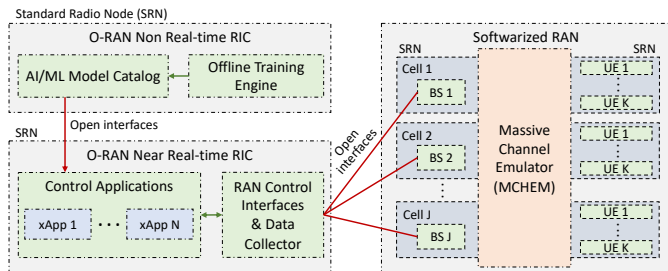


Fig. 6: Example of O-RAN integration in Colosseum.

Based on the developments of [26], we now showcase how Colosseum can be used to instantiate an O-RAN-managed cellular infrastructure. The performance of the network is optimized through data-driven closed-loop control enabled by xApps deployed at the near-real time RIC. We deploy 4 softwarized base stations and 40 users in a dense urban scenario. The base stations divide the available spectrum in three slices—Enhanced Mobile Broadband (eMBB), Machine-type

Communications (MTC), and Ultra Reliable and Low Latency Communication (URLLC)—and they run our SCOPE framework [25], which enables dynamic control of slicing allocation and scheduling policies. We control the configuration of each slice in real time through Deep Reinforcement Learning (DRL) agents deployed as xApps on the near-real time RIC. The agents receive periodic data reports from the base stations and send them control directives through the O-RAN E2 termination. An example of these control directives is the selection of which scheduling policy among Round-robin (RR), Waterfilling (WF), and Proportionally Fair (PF) to implement for each network slice of each base stations.




Fig. 7: Resource allocation ratio of the URLLC slice for different scheduling policies and with DRL control.

Figure 7 shows the Cumulative Distribution Function (CDF) of the resources allocated to the URLLC slice when the RR, WF, and PF scheduling policies are implemented statically, and with DRL control. For this slice, the agents aim at maximizing the ratio of resources granted to the users and those they request. Since the DRL agents dynamically choose the optimal scheduling policy based on the current channel conditions and traffic demand, they achieve superior performance compared to the static scheduling policies.

## B. Wi-Fi

Wi-Fi has become an indispensable and ubiquitous wireless technology to provide home networking and public/hotspot Internet connectivity, as well as supporting a wide range of Internet of Things (IoT) applications. To support the ever-increasing number of devices, the Federal Communications Commission (FCC) has recently opened an additional 1.2 GHz of spectrum in the 6 GHz band [35]. This portion of the spectrum, intended for unlicensed use, has been embraced in the IEEE 802.11ax amendment, namely Wi-Fi 6 [36]. Large-scale experimentation is a core enabler for research and development in these emerging spectrum bandwidths.

To showcase how Colosseum can facilitate these operations, we instantiate a Wi-Fi network through an open-source GNU Radio-based implementation of the IEEE 802.11a/g/p standard [37]. We deploy 50 Wi-Fi transceivers in a Colosseum urban scenario and measure the Signal to Interference plus Noise Ratio (SINR) in the case nodes do not move (Fig. 8, top), and in the case they move with an average speed of 3 m/s (Fig. 8, bottom). We notice that when node mobility




Fig. 8: SINR measured by Wi-Fi nodes on Colosseum in a scenario without node mobility (top) and with node mobility (bottom).

is implemented, the SINR degrades by 3.18 dB on average because of the rapidly-changing channel conditions.

### C. Spectrum Sharing

In recent years, the widespread use of wireless communications has fostered technological advancement and motivated the development and deployment of innovative services for mobile users. These phenomena have resulted in a sharp increase in the number of wireless devices and traffic demand. This has exacerbated the issue of the lack of spectrum to support them, also known as *spectrum crunch* [2, 38].

In most cases, spectrum is either allocated to government agencies for research and military uses (e.g., radars, law enforcement), to the scientific community (e.g., for radio astronomy, atmospheric sensing), or licensed to operators, which gain exclusive access to spectrum frequencies in exchange for very high licensing fees. Only in a handful of cases, such as the Industrial, Scientific and Medical (ISM) band, spectrum is allocated to non-commercial and experimental applications. One of the main drawbacks of such a static allocation is that unlicensed portions of the spectrum are extremely congested, while licensed ones might experience lower load conditions due to the lack of activity.

In this context, one research topic that has gained momentum is spectrum sharing, whose ultimate goal is to devise solutions that allow multiple technologies to share portions of the spectrum. To this end, SDRs have gained increasing attention in the community as a promising solution to allow unlicensed users to fill the so-called *spectrum holes* left by licensed users. Although the literature is characterized by several works that leverage optimization and data-driven solutions in SDRs, one of the main limitations of such works is the lack of large-scale experimentation that demonstrates their effectiveness in heterogeneous network topologies, traffic and channel conditions [7, 25].

Colosseum fills this gap and allows users to instantiate large-scale networks with nodes running heterogeneous wireless protocol stacks (e.g., cellular, Wi-Fi) on the same spectrum bandwidth. This enables researchers to design, implement and evaluate novel spectrum sharing solutions on heterogeneous and diverse network deployments and conditions without caus-

ing harmful interference to licensed users. To provide an example of this application, we deploy on the same portion of the spectrum a cellular base station serving 5 cellular users, and two Wi-Fi nodes. Figure 9 shows the average Channel Quality Information (CQI) and percentage of uplink errors experienced by the cellular users with and without Wi-Fi transmissions active. We notice that the interference generated by Wi-Fi




Fig. 9: Average cell CQI and percentage of uplink errors with and without Wi-Fi traffic. Time periods without Wi-Fi traffic are marked with a red shaded area.

communications causes the CQI of the cellular users to rapidly drop (Fig. 9a), and the percentage of uplink errors to increase (Fig. 9b) when Wi-Fi transmissions are ongoing.

### D. Unmanned Aerial Vehicles

The growing commercialization of UAVs, along with their greater affordability and increasing popularity, has recently spawned the interest of telecom operators and equipment providers toward the use of drones for networking applications. Examples of such applications include the use of UAVs to create swarm of flying networks [3, 39], to promptly aid in the aftermath of disaster scenarios [40–42], and to bring connectivity to remote areas [43, 44]. However, security concerns, strict flying regulations, and the cost of the equipment itself impose additional challenges to the development and testing of solutions at scale.

Colosseum has the potential of facilitating the evaluation of such solutions by providing the means to emulate large-scale SDR-enabled UAV networks. To showcase these capabilities, we instantiate a network with 30 nodes in the Alleys of Austin scenario (see Section III-A). Nodes are divided in three groups, or *squads* (see Fig. 10a). Each squad is composed of a UAV (marked with “U” in the figure) and 9 ground relay nodes (“R”), which need to deliver data (e.g., video and voice traffic) to specific members of the squad.

At each instant of time, specific nodes act as relays (aerial or ground) for data from a squad to the intended destination (e.g., nodes in other squads), while coping with inter-user and inter-squad interference. Figure 10b shows the average SINR when two different algorithms are used to select the relay nodes:

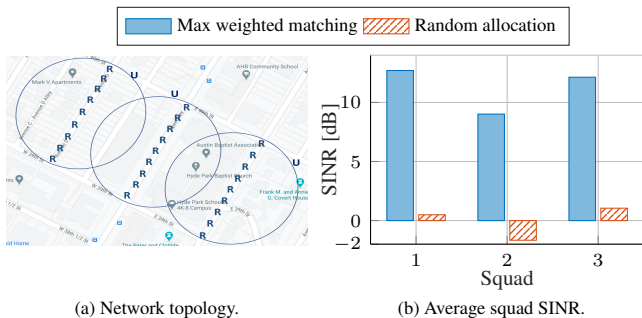


Fig. 10: Left: network topology of an experiment with 30 nodes (ground relays, “R”, and UAVs, “U”) in the Alleys of Austin scenario. The circles mark the trajectory of the UAVs. Right: average SINR of three squads in Alleys of Austin scenario when different algorithms are used for the selection of relay nodes.

SINR-based maximum weighted matching, and random allocation. We notice that when the maximum weighted matching algorithm is used, the squads experience a higher SINR than when the random allocation strategy is enforced. This is due to the fact that the former algorithm elects as relays those nodes that have better channel conditions (and thus a higher SINR) toward the remaining squad members. This algorithm also takes advantage of the UAVs as aerial relays, which are in line-of-sight condition with the ground nodes.

## VI. EVOLUTION OF COLOSSEUM

While Colosseum represents an invaluable tool for the evaluation of large-scale wireless networking scenarios, it also needs to evolve to match the needs of communication technologies in the years to come. The following sections describe planned extensions of Colosseum that will add new emulation capabilities (Section VI-A) and AI components (Section VI-B).

### A. Expansions to Support NRDZ and mmWave Research

A steady trend in the wireless industry has been the push for the allocation of additional spectrum for communications, with spectrum sharing [45] and the adoption of new frequency bands, e.g., millimeter waves [46, 47]. As the usage of the spectrum expands into new frontiers, it becomes important to study and understand how different spectrum users can safely coexist. Along this line, the NSF has recently published a “Dear Colleague Letter” to explore the feasibility of establishing National Radio Dynamic Zones (NRDZs) in the United States. NRDZs are meant as safe platforms for experiments that would not be allowed under FCC spectrum regulations because of chances to generate harmful interference toward incumbent users. Colosseum is naturally positioned to implement an NRDZ, as—by being an emulator—it avoids any interference to external devices. Additionally, its programmability allows experimenters to test multiple scenarios, different protocol stacks, and different traffic patterns (see Section III).

As part of the NRDZ activities, Colosseum will be expanded to support new spectrum bands, scenarios, and power levels. Colosseum already supports very granular coordinated spectrum sharing across heterogeneous networks, and will act as a “management and optimizer” system for protection, observation, validation and automation of spectrum sharing. These




Fig. 11: Current components of Colosseum and proposed extensions for the NRDZ program.

capabilities will be extended by adding two new quadrants (see Fig. 11) and by improving the current MCHM emulation.

As of today, there are two limitations in Colosseum that prevent it to act as a tool to design and study NRDZs as defined above. First, Colosseum was designed to emulate terrains up to  $1 \times 1$  km. This limitation is a consequence of the maximum propagation delay that can be emulated by MCHM, which is embedded in the design of its FPGAs. Second, Colosseum was designed to emulate sub-6 GHz omnidirectional transmissions. However, it cannot emulate directional transmissions and higher frequency bands.

To go beyond today’s limitations, we are deploying an additional, reduced version of a Colosseum quadrant as a development environment. This environment includes a dedicated MCHM quadrant that will be the basis for the development of new channel emulation code, as well as the testing of new hardware for the user radios and SRNs. Notably, two extensions for MCHM will be introduced first in this environment, and then scaled to the full Colosseum (Fig. 11). The first will allow longer delays for high-fidelity emulation of a  $100 \times 100$  km grid and terrains representative of NRDZs. The second will introduce the modeling of beamforming-based, directional communications to emulate mmWave networks. Finally, to support the bandwidth requirements of 5G and beyond cellular networks, we will create a new Colosseum quadrant able to emulate very large baseband bandwidth (up to 2 GHz of bandwidth).

### B. AI Jumpstart Integration

Colosseum is currently being expanded through the MassTech “AI Jumpstart” program [48]. This program aims at jump-starting Massachusetts firms interested in deploying AI to enhance their businesses, connecting industry practitioners with world-class facilities and researchers. Part of the AI Jumpstart equipment will be integrated in the Colosseum environment to enable research and development of AI-based wireless networking solutions. Examples of use cases are the fast and efficient training of large-scale wireless datasets collected on Colosseum [26]; real-time, AI-driven 5G signal processing for the full 5G protocol stack [49]; and model-free adaptation and control of large-scale wireless networks.

The new equipment includes two NVIDIA DGX A100 nodes—among the most powerful AI compute solutions on the market today—capable of delivering up to 10petaFLOPS [50]. These nodes include 8 GPUs, state-of-the-art AMD CPUs, 1 TB of RAM, and 10 Mellanox ConnectX-6 network cards, each capable of sustaining a 200 Gbps link. An additional large

memory node enables memory-intensive workloads, with 3 TB of RAM. These machines are connected through a dedicated Mellanox Infiniband switch, which can sustain an aggregated traffic of up to 16 Tbps. This system will be fully meshed with the compute and wireless resources of Colosseum, with dedicated Nomad-based orchestration and load balancing capabilities [51].

## VII. CONCLUSIONS

This paper presents Colosseum—the world’s largest wireless network emulator with hardware-in-the-loop—as a testbed that is *for the first time* publicly available to the community. Colosseum enables experimental research at scale through virtualized and softwarized waveforms and protocol stacks. The unique emulation capabilities of Colosseum enable the design, prototyping, and testing of wireless solutions in a host of scenarios and channel conditions. We illustrate the architecture, emulation capabilities and operational modes of Colosseum. We then provide examples of large-scale experimentation in Colosseum, considering cellular, Wi-Fi, spectrum sharing and UAV scenarios. We finally lay out a plan for extending Colosseum to enable experiments over larger deployments and higher frequencies.

## REFERENCES

- [1] 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. 1–18, December 2020.
- [2] G. Naik and J.-M. J. Park, “Coexistence of Wi-Fi 6E and 5G NR-U: Can We Do Better in the 6 GHz Bands?” in *Proceedings of IEEE INFOCOM*, Vancouver, BC, Canada, May 2021.
- [3] L. Bertizzolo *et al.*, “SwarmControl: An Automated Distributed Control Framework for Self-Optimizing Drone Networks,” in *Proceedings of IEEE INFOCOM*, Toronto, ON, Canada, July 2020.
- [4] J. Buczek, L. Bertizzolo, S. Basagni, and T. Melodia, “What is A Wireless UAV? a design blueprint for 6G flying wireless nodes,” in *Proceedings of ACM WiNTECH 2021*, New Orleans, LA, January 2022.
- [5] G. E. Santagati and T. Melodia, “A Software-Defined Ultrasonic Networking Framework for Wearable Devices,” *IEEE/ACM Transactions on Networking*, vol. 25, no. 2, pp. 960–973, April 2017.
- [6] J. Breen *et al.*, “POWDER: Platform for Open Wireless Data-driven Experimental Research,” *Computer Networks*, vol. 197, pp. 1–18, October 2021.
- [7] R. Doost-Mohammady, O. Bejarano, and A. Sabharwal, “Good Times for Wireless Research,” *Computer Networks*, vol. 188, pp. 1–9, April 2021.
- [8] M. Kohli *et al.*, “Open-Access Full-Duplex Wireless in the ORBIT and COSMOS Testbeds,” *Computer Networks*, 2021.
- [9] A. Panicker *et al.*, “AERPAW Emulation Overview and Preliminary Performance Evaluation,” *Computer Networks*, vol. 194, pp. 1–11, July 2021.
- [10] H. Zhang *et al.*, “ARA: A Wireless Living Lab Vision for Smart and Connected Rural Communities,” in *Proceedings of ACM WiNTECH*, New Orleans, LA, USA, October 2021.
- [11] H. Koumaras *et al.*, “5GENESIS: The Genesis of a Flexible 5G Facility,” in *Proceedings of IEEE CAMAD*, Barcelona, Spain, September 2018.
- [12] K. R. Dandekar *et al.*, “Grid Software Defined Radio Network Testbed for Hybrid Measurement and Emulation,” in *Proceedings of IEEE SECON*, Boston, MA, USA, June 2019.
- [13] L. Bertizzolo *et al.*, “Arena: A 64-antenna SDR-based Ceiling Grid Testing Platform for Sub-6 GHz 5G-and-Beyond Radio Spectrum Research,” *Computer Networks*, pp. 1–17, November 2020.
- [14] DARPA. Spectrum Collaboration Challenge. <https://www.darpa.mil/program/spectrum-collaboration-challenge>. Accessed September 2021.
- [15] P. Tilghman, “AI Will Rule the Airwaves: A DARPA Grand Challenge Seeks Autonomous Radios to Manage the Wireless Spectrum,” *IEEE Spectrum*, vol. 56, no. 6, pp. 28–33, May 2019.
- [16] R. L. Yuan and K. M. Schmidt, “Defense Advanced Research Projects Agency Spectrum Collaboration Challenge at APL: Introduction,” vol. 35, no. 1, pp. 2–3, 2019.
- [17] D. Coleman *et al.*, “Overview of the Colosseum: The World’s Largest Test Bed for Radio Experiments,” *Johns Hopkins APL Technical Digest*, vol. 35, no. 1, pp. 4–11, 2019.
- [18] A. S. Freeman *et al.*, “Software Project Management for the Defense Advanced Research Projects Agency Spectrum Collaboration Challenge,” *Johns Hopkins APL Technical Digest*, vol. 35, no. 1, pp. 12–21, 2019.
- [19] A. T. Plummer, Jr. and K. P. Taylor, “Development and Operations on the Defense Advanced Research Project Agency’s Spectrum Collaboration Challenge,” *Johns Hopkins APL Technical Digest*, vol. 35, no. 1, pp. 22–33, 2019.
- [20] J. W. Mok, A. L. Hom, J. J. Uher, and D. M. Coleman, “The Resource Manager for the Defense Advanced Research Projects Agency Spectrum Collaboration Challenge Test Bed,” *Johns Hopkins APL Technical Digest*, vol. 35, no. 1, pp. 34–41, 2019.
- [21] D. A. White, Jr., J. E. Annis, and F. F. Johnson, “Standard Radio Nodes in the Defense Advanced Research Projects Agency Spectrum Collaboration Challenge,” vol. 35, no. 1, pp. 42–48, 2019.
- [22] K. J. Yim, K. R. McKeever, and D. R. Barcklow, “Incumbent Radio Systems in the Defense Advanced Research Projects Agency Spectrum Collaboration Challenge Test Bed,” vol. 35, no. 1, pp. 49–57, 2019.
- [23] P. D. Curtis, A. T. Plummer, Jr., J. E. Annis, and W. J. La Cholter, “Traffic Generation System for the Defense Advanced Research Projects Agency Spectrum Collaboration Challenge,” *Johns Hopkins APL Technical Digest*, vol. 35, no. 1, pp. 58–68, 2019.
- [24] D. Barcklow *et al.*, “Radio Frequency Emulation System for the Defense Advanced Research Projects Agency Spectrum Collaboration Challenge,” *Johns Hopkins APL Technical Digest*, vol. 35, no. 1, pp. 69–78, 2019.
- [25] L. Bonati, S. D’Oro, S. Basagni, and T. Melodia, “SCOPE: An Open and Softwarized Prototyping Platform for NextG Systems,” in *Proceedings of ACM MobiSys*, Virtual Conference, June 2021.
- [26] 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*, 2021.
- [27] A. Chaudhari and M. Braun, “A Scalable FPGA Architecture for Flexible, Large-Scale, Real-Time RF Channel Emulation,” in *Proceedings of IEEE ReCoSoC*, Lille, France, July 2018.
- [28] U.S. Naval Research Laboratory. MGEN Traffic Emulator. <https://www.nrl.navy.mil/Our-Work/Areas-of-Research/Information-Technology/NCS/MGEN>. Accessed September 2021.
- [29] M. Polese, L. Bertizzolo, L. Bonati, A. Gosain, and T. Melodia, “An Experimental mmWave Channel Model for UAV-to-UAV Communications,” in *Proceedings of ACM mmNets*, London, United Kingdom, September 2020.
- [30] M. Tehrani-Moayyed, L. Bonati, P. Johari, T. Melodia, and S. Basagni, “Creating RF Scenarios for Large-Scale, Real-Time Wireless Channel Emulators,” in *Proceedings of IEEE MedComNet*, Ibiza, Spain, June 2021.
- [31] Unwired Labs, “OpenCelliD,” accessed September 2021. [Online]. Available: <https://opencellid.org>
- [32] S. D’Oro, L. Bonati, F. Restuccia, and T. Melodia, “Coordinated 5G Network Slicing: How Constructive Interference Can Boost Network Throughput,” *IEEE/ACM Transactions on Networking*, vol. 29, no. 4, pp. 1881–1894, August 2021.
- [33] L. Bonati *et al.*, “CellOS: Zero-touch Softwarized Open Cellular Networks,” *Computer Networks*, vol. 180, pp. 1–13, October 2020.
- [34] I. Gomez-Miguel *et al.*, “srsLTE: An Open-source Platform for LTE Evolution and Experimentation,” in *Proceedings of ACM WiNTECH*, New York City, NY, USA, October 2016.
- [35] FCC. (2020, April) FCC Adopts New Rules for the 6 GHz Band, Unleashing 1,200 MHz of Spectrum for Unlicensed Use. <https://docs.fcc.gov/public/attachments/DOC-363945A1.pdf>. Accessed September 2021.
- [36] IEEE 802.11 Wireless LAN Working Group, “IEEE Standard for Information Technology–Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks–Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 1: Enhancements for High-Efficiency WLAN,” *IEEE Std 802.11ax-2021 (Amendment to IEEE Std 802.11-2020)*, pp. 1–767, May 2021.
- [37] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, “Performance Assessment of IEEE 802.11p with an Open Source SDR-based Prototype,” *IEEE Transactions on Mobile Computing*, vol. 17, no. 5, pp. 1162–1175, May 2018.



- [38] P. Yang, L. Kong, and G. Chen, "Spectrum Sharing for 5G/6G URLLC: Research Frontiers and Standards," *IEEE Communications Standards Magazine*, vol. 5, no. 2, pp. 120–125, June 2021.
- [39] S. Mohanti *et al.*, "AirBeam: Experimental Demonstration of Distributed Beamforming by a Swarm of UAVs," in *Proceedings of IEEE MASS*, Monterey, CA, USA, November 2019.
- [40] A. Coletta, G. Maselli, M. Piva, and D. Silvestri, "DANGER: A Drones Aided Network for Guiding Emergency and Rescue Operations," in *Proceedings of ACM Mobihoc*, Virtual Conference, October 2020.
- [41] L. Ferranti, S. D'Oro, L. Bonati, F. Cuomo, and T. Melodia, "HIRO-NET: Heterogeneous Intelligent RObotic Network for Internet sharing in Disaster Scenarios," *IEEE Transactions on Mobile Computing*, 2021.
- [42] C. Rottondi, F. Malandrino, A. Bianco, C. F. Chiasserini, and I. Stavrakakis, "Scheduling of Emergency Tasks for Multiservice UAVs in Post-disaster Scenarios," *Computer Networks*, vol. 184, pp. 1–13, January 2021.
- [43] L. Bertizzolo *et al.*, "mmBAC: Location-aided mmWave Backhaul Management for UAV-based Aerial Cells," in *Proceedings of ACM mmNets*, Los Cabos, Mexico, October 2019.
- [44] L. Ferranti, L. Bonati, S. D'Oro, and T. Melodia, "SkyCell: A Prototyping Platform for 5G Aerial Base Stations," in *Proceedings of IEEE SwarmNet*, Cork, Ireland, August 2020.
- [45] L. Zhang *et al.*, "A Survey of Advanced Techniques for Spectrum Sharing in 5G Networks," *IEEE Wireless Communications*, vol. 24, no. 5, pp. 44–51, October 2017.
- [46] S. Rangan, T. S. Rappaport, and E. Erkip, "Millimeter-Wave Cellular Wireless Networks: Potentials and Challenges," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 366–385, March 2014.
- [47] M. Giordani, M. Polese, A. Roy, D. Castor, and M. Zorzi, "A Tutorial on Beam Management for 3GPP NR at mmWave Frequencies," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 173–196, February 2019.
- [48] The Innovation Institute at the MassTech Collaborative. (2020, February) AI Jumpstart. <https://innovation.masstech.org/projects-and-initiatives/collaborative-research-matching-grant-program/ai-jumpstart>. Accessed September 2021.
- [49] M. Polese, F. Restuccia, and T. Melodia, "DeepBeam: Deep Waveform Learning for Coordination-Free Beam Management in mmWave Networks," in *Proceedings of ACM Mobihoc*, Shanghai, China, July 2021.
- [50] NVIDIA. NVIDIA DGX A100—The Universal System for AI Infrastructure. <https://www.nvidia.com/en-us/data-center/dgx-a100>. Accessed September 2021.
- [51] HashiCorp. Nomad. <https://www.nomadproject.io>. Accessed September 2021.