Computer Networks 181 (2020) 107436

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

# Arena: A 64-antenna SDR-based ceiling grid testing platform for sub-6 GHz 5G-and-Beyond radio spectrum research

Lorenzo Bertizzolo [*,a], Leonardo Bonati [a], Emrecan Demirors [a], Amani Al-shawabka [a], Salvatore D'Oro [a], Francesco Restuccia [a], Tommaso Melodia [a]

*Institute for the Wireless Internet of Things, Northeastern University, Boston, MA 02115, USA*

ARTICLE INFO

*Keywords:*
Software-defined radios
Wireless testbed
Spectrum research
Antenna grid
Internet of things

ABSTRACT

Arena is an open-access wireless testing platform based on a grid of antennas mounted on the ceiling of a large office-space environment. Each antenna is connected to programmable software-defined radios (SDR) enabling sub-6 GHz 5G-and-beyond spectrum research. With 12 computational servers, 24 SDRs synchronized at the symbol level, and a total of 64 antennas, Arena provides the computational power and the scale to foster new technology development in some of the most crowded spectrum bands. Arena is based on a three-tier design, where the servers and the SDRs are housed in a double rack in a dedicated room, while the antennas are hung off the ceiling of a 2240 square feet office space and cabled to the radios through 100 ft-long cables. This ensures a reconfigurable, scalable, and repeatable real-time experimental evaluation in a real wireless indoor environment.

In this paper, we introduce the architecture, capabilities, and system design choices of Arena, and provides details of the software and hardware implementation of various testbed components. Furthermore, we describe key capabilities by providing examples of published work that employed Arena for applications as diverse as synchronized MIMO transmission schemes, multi-hop ad hoc networking, multi-cell 5G networks, AI-powered Radio-Frequency fingerprinting, secure wireless communications, and spectrum sensing for cognitive radio.

## 1. Introduction

The evolution of wireless networked systems continues to be a crucial commercial, strategic, and geopolitical matter. According to the latest Ericsson mobility report, there are now 5.7 billion mobile broadband subscriptions worldwide, generating more than 130 exabytes per month of wireless traffic [1]. Moreover, it is expected that by 2020, over 50 billion devices will be absorbed into the Internet, generating a global network of "things" of dimensions never seen before—and growing at a compound annual growth rate (CAGR) of 27% [2]. Industrial automation, smart cities, and distributed robotic systems will increasingly rely on large-scale wireless networked systems. In addition to their commercial strategic need, 5G and IoT have been identified as critical technologies for national security. Thus, wireless systems will continue to change the way we live, work, manufacture goods, and provide national security.

It is therefore to some extent surprising that the wireless research community is still lacking experimental facilities to support a "science" of rigorous and repeatable experimental evaluation of wireless networked systems—beyond simulation tools and small-scale, ad hoc, testbeds. The recent NSF Platforms for Advanced Wireless Research (PAWR) program [3] is attempting to address this by developing four city-scale platforms for advanced wireless research to experiment with new IoT and wireless systems in outdoor "out-in-the-wild" environments [4]. Similarly, Colosseum [5], transitioned from the DARPA Spectrum Collaboration Challenge to Northeastern University [6], will provide a *shared* wireless network emulation facility to experiment and test *at scale*, in a fully *controlled and observable environment*, and with *hardware in the loop*. While the availability of PAWR platforms and Colosseum is expected to be a major stepping stone toward the goal of open rigorous experimentation with shared facilities, the community is still lacking a platform to test at scale medium- and short-range radio technologies in the sub-6 GHz radio bands in an indoor realistic environment able to guarantee high-fidelity, scale, and repeatability of experiments. This is crucial for sub-6 GHz testing indoor deployments such as offices, malls, and airports that are characterized by fast-varying environment, spatially and time-varying interference, significant multi-path effect, and continuous mobility of surrounding objects.

* Corresponding author.
*E-mail address:* bertizzolo@ece.neu.edu (L. Bertizzolo).

https://doi.org/10.1016/j.comnet.2020.107436
Received 14 April 2020; Received in revised form 2 July 2020; Accepted 17 July 2020
Available online 28 July 2020
1389-1286/

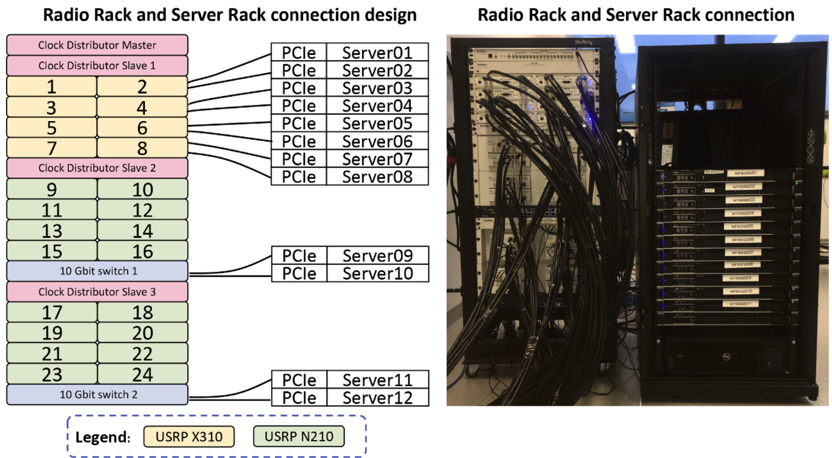


**Fig. 1.** Arena's Server and Radio Rack. The 24 radios are located in a single room together with the servers driving them. This design choice allows easy synchronization across the radios while the antennas are installed in another space and connected through long coaxial cables.

To address this need, in [7,8] we introduced for the first time *Arena*, an open-access wireless testing platform based on an indoor 64-antenna ceiling grid, connected to programmable SDRs for sub-6 GHz 5G-and--beyond spectrum research. Arena is located in the open-space laboratory on the fourth floor of the Northeastern University Interdisciplinary Science & Engineering Complex. The key characteristics of Arena can be summarized as follows:

- **Real-time evaluation platform.**Arena is an open-access testing platform to run experiments in real-time with SDRs. It is accessible through the Internet and it can be used to prototype, develop, and experimentally evaluate new emerging wireless and edge computing technologies.
- **Over-the-air evaluation.**Arena consists of 64 antennas covering an area of 2240 ft$^2$ for a total of 4032 real wireless channels, with distances among each pair of antennas ranging from 5 ft to 85 ft. Arena's coverage area, indoor line-of-sight distances, and scale are unprecedented for an office deployment. It facilitates rigorous and extensive testing of wireless technologies on real, uncontrollable, over-the-air wireless channels in a real, high-fidelity indoor office environment.
- **Fully-synchronized radios.**Arena is based on a 24-SDR rack driving a total of 64 transmit/receive antennas deployed on a ceiling grid layout. The radios are synchronized via clock distributors and connected to the antennas using identical equal-length cables, ensuring full symbol-level synchronization throughout the whole testbed. This enables applications such as massive MIMO, cooperative multi-point MIMO, and synchronized distributed systems.
- **Repeatable, flexible, and scalable indoor experiments.**Arena's 64 antenna grid provides a plethora of possible network topologies and the scale to foster new technology development. Arena's design ensures unchanged locations of the antennas throughout the experiments and guarantees the integrity of the collected experimental data. A second server rack contains 12 identical servers that control the radios and perform the baseband processing operations. Arena's three-tier design made of the server rack, the radio rack, and the ceiling grid addresses typical SDR deployment issues such as antenna orientation, cable non-linearities, and, ultimately, guarantees repeatability of experiments.

*In a way, Arena — with its scale, nodes' distances, and real wireless channel in a contiguous, uncontrolled, and wireless-diverse office environment characterized by multiple surrounding objects, rich multi-path, and unpredictable people mobility — represents an ideal testing platform for WLAN and indoor cellular technologies.* In this article, we review in detail the design choices, the hardware and software components of Arena, and showcase several application scenarios that leverage Arena for performance evaluation of wireless systems. The rest of this paper is organized as follows.

In Section 2, we outline the testbed design and system architecture, while in Section 3 we describe Arena's hardware and software components. In Section 4, we discuss how to access Arena, while in Section 5 we showcase a series of applications scenarios. Finally, related work is surveyed in Section 6, while we draw our main conclusions in Section 7.

## 2. Testing platform design and system architecture overview

The motivation behind this effort is the need for rigorous experimental evaluation in an indoor environment, with real, uncontrollable,

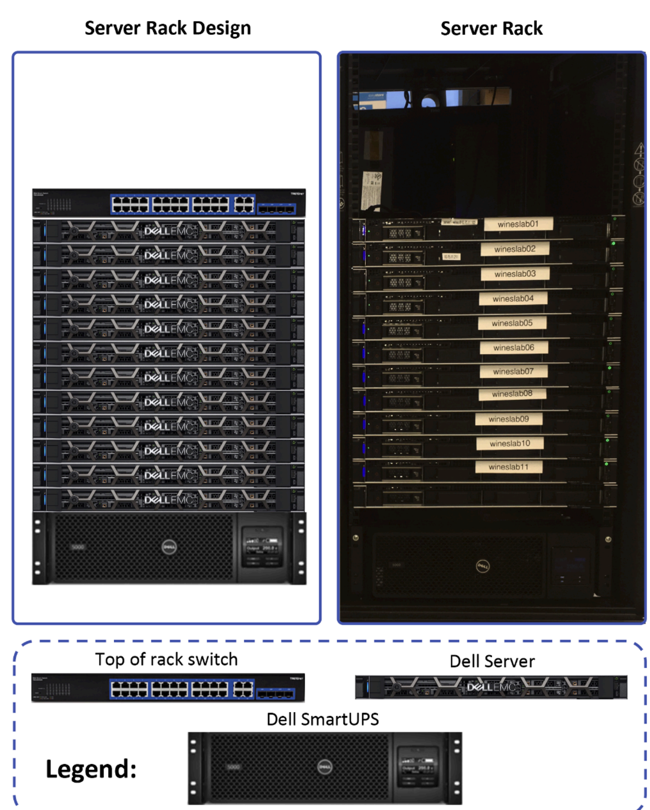


**Fig. 2.** Arena's Server Rack design and installation.

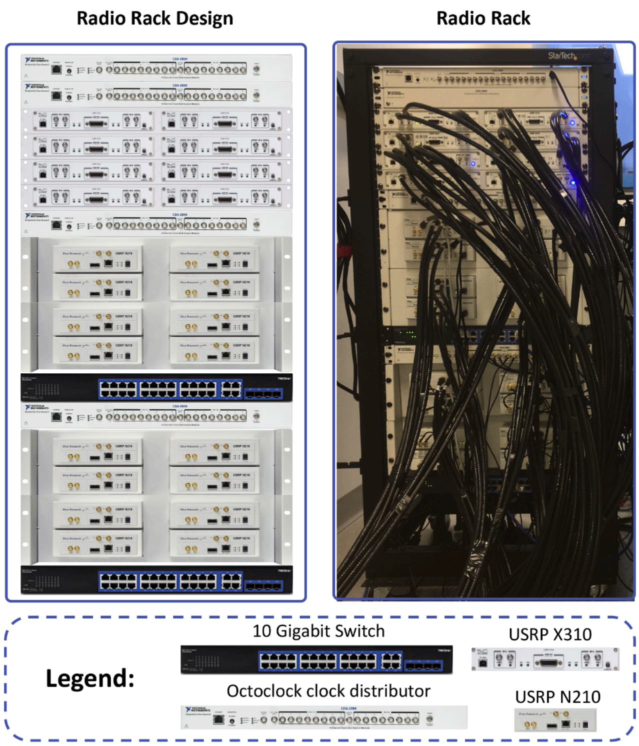


**Fig. 3.** Arena's Radio Rack design and installation.

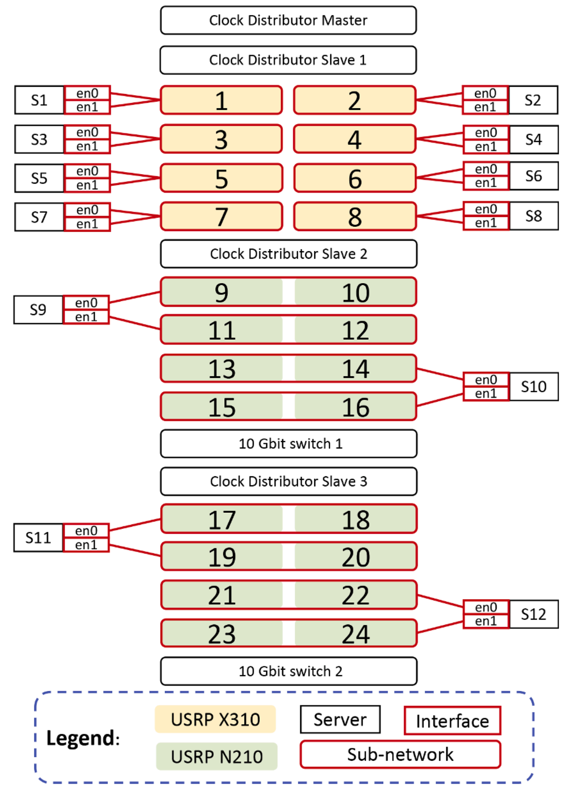


**Fig. 4.** Arena's networking configuration. All servers S1–S12 are equipped with a PCIe network card offering two additional 10 Gigabit interfaces `en0` and `en1`. S1—S8 drive one USRP X310 using the two interfaces combined, while each serves S9—S12's interface drive two USRP N210 as shown in red color. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

wireless channels, highly representative of the final indoor deployment scenario. Arena provides researchers with a software control framework and radio hardware to evaluate wireless development on a multitude of different radio configurations, topologies, and channel conditions. More importantly, it offers the capability of scaling up the testing environment in software without requiring any physical radio relocation. Arena is based on a three-tiered architecture: the server rack, the radio rack, and the antenna grid, as illustrated in Fig. 1.

### 2.1. The server rack.

The server rack consists of 12 servers individually accessible through a top-of-the-rack gateway responsible for authenticating users, as illustrated in Fig. 2. The Gateway is in charge of granting access to authenticated users only, and it interfaces with the university intranet via a gigabit Ethernet connection. The platform is thus accessible from both the Northeastern University intranet and the Internet via the College of Engineering (COE) Gateway, which keeps an up-to-date list of allowed users.

The Server Rack is in charge of driving the radios and of performing all the baseband processing operations. The Server Rack hosts twelve identical servers both on hardware and software capabilities. They feature the same kernel version, operating system, and installed software, to guarantee uniform computational power and fair operations across the whole testbed. Upon access, the user's disk space is mounted on the servers through the Network File System (NFS) protocol so as to guarantee files and new software consistency across the whole Server Rack.

As mentioned, the servers drive the radios and perform the computational-heavy baseband processing of transmission and reception operations. The servers connect to the radios through a dedicated PCIExpress network card, which offers two additional network interfaces at 10 Gbit/s each. Each server is also connected to the gateway through a standard 1 Gbit/s interface. With the goal of load balancing the computational-intensive baseband processing operations, each server drives one or more radios in one of two possible configurations: *1-to-1* driving, and *1-to-4* driving.

*1-to-1 driving*: Eight of the 12 servers are connected one-to-one to individual Universal Software Radio Peripheral (USRP) X310 SDRs via a dedicated 10 Gigabit network interface. The USRP X310 is a high-performance, scalable SDR platform whose hardware architecture combines two extended-bandwidth daughterboard slots covering DC–6 GHz with up to 120 MHz of baseband bandwidth. The USRP X310 architecture includes two 10 Gigabit Ethernet interfaces to stream data to and from host processors. The wide maximum bandwidth of these devices asks for significant baseband processing power to operate in real-time, therefore each USRP X310 can rely on a dedicated server. The eight servers and the 8 USRP X310 couples are connected through a 10 SFP+ Gigabit interface that provides high bandwidth and a low-latency connection between servers and SDRs.

*1-to-4 driving*: The remaining four servers drive four Universal Software Radio Peripheral (USRP) N210 SDRs each, for a total of 16 USRPs N210. The USRP N210 is a high-bandwidth, high-dynamic range radio designed to operate from DC to 6 GHz with up to 56 MHz of baseband bandwidth. The USRP N210 architecture includes a Gigabit Ethernet connectivity to stream data to and from host processors. To connect the control hosts to the SDRs, we employ two dedicated 10 Gigabit switch. Each 10 Gigabit switch offers multiple speed connections and connects the servers over 10 Gigabit interfaces to the eight USRP N210 over 1 Gigabit Ethernet interfaces.

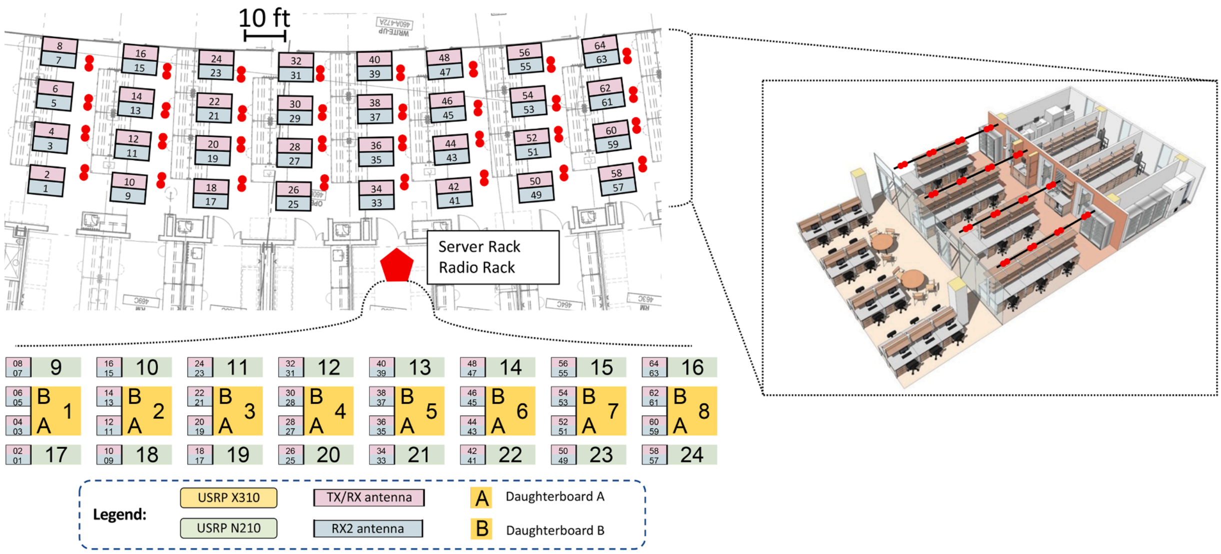


**Fig. 5.** Arena's Antenna Grid layout. The USRP X310, which house two daughterboards, are in the middle and stretch along the whole length of the deployment. These SDRs are more powerful, control 4 antennas total, and can implement LTE-compliant base stations (eNB/gNB) and users (UEs). USRP N210 stretch along the whole deployment length forming a 8 × 2 -SDR symmetrical grid.

### *2.2. The radio rack.*

As mentioned in the previous section, the Radio Rack is composed of sixteen USRPs N210, eight USRPs X310, four clock distributors, and two 10 Gigabit switches, for a total of 24 synchronized SDRs. These house 32 daughterboards total, each of which hosting one TX/RX and one RX2 antenna terminals. The Radio Rack layout is shown in Fig. 3. From a synchronization perspective, the SDRs in the rack are logically organized into 8-radio groups, each of which is provided with time and frequency synchronization by a dedicated OctoClock clock distributor. The three clock distributors are in turn synchronized to a master clock distributor, which can generate time and frequency signals. To connect the master to the three slave clock distributors and the latter to the 24 SDRs, we employed 54 identical and same-length cables. Identical cables guarantee clear reference signals and identical delays across all the 24 radios, which is essential for full synchronization across the Radio Rack. From a synchronization accuracy standpoint, Arena represents a one of a kind platform where the antenna terminals are deployed in a wide indoor space, while the radio boards are co-located and centrally synchronized with clock distributor precision (25 ppb), thus avoiding the inaccuracy of distributed synchronization mechanisms. Further, the RF daughterboards have been individually calibrated disconnecting any external hardware as instructed in [9] as so to minimize the RX IQ imbalance vs. LO frequency, the TX DC offset vs. LO frequency, and the TX IQ imbalance vs. LO frequency. The calibration files have been made available in all Arena servers and are automatically loaded when calling the UDH SDR driver. Last, even though the radios are externally synchronized in phase and frequency, and the daughterboards have been calibrated, the users are nevertheless invited to measure and possibly compensate in software the residual constant phase offset (see [10] for more details) as discussed in [11]. We will see later on that ensuring time and frequency synchronization is a fundamental enabler of applications such as MIMO communication schemes and cooperative multi-point transmissions.

From a networking perspective, the SDRs and the control hosts have been configured into separate sub-networks with the goal of load balancing the traffic across the control host interfaces. As anticipated, each USRP X310 is on a private sub-network with its only control host (1-to-1 driving), while USRPs N210 are split into 4-SDRs groups driven by one server each (1-to-4 driving). Specifically, each 4-SDR group is configured into two sub-networks, each of which is reachable from one of the server's 10 Gigabit interfaces only as illustrated in Fig. 4.

This network configuration breaks down the traffic over the two server interfaces and makes sure that no interface drives more than 2 USRP N210 or one USRP X310. Given the baseband processing capabilities of USRPs X310 (200 MSamples/s at 32 bit/Sample) tops) and of USRPs N210 (25 MSamples/s at 32 bit/Sample) tops), this configuration aims at load balancing the traffic across the servers and their two 10 Gigabit/s network interfaces with the ultimate guarantee of correct operations even at full bandwidth capacity.

### *2.3. The antenna grid*

As previously mentioned, each SDR houses one (USRP N210) or two (USRP X310) radio frequency daughterboards, each of which hosts two antenna terminals, namely TX/RX and RX2. Each terminal is individually wired to one single antenna hanging off the ceiling of on open-space laboratory through a 100 ft long coaxial cable. Antenna terminals of the same daughterboard connect to one *antenna pair*, which is a two-antenna group hanging off the ceiling approximately at the same location. Each USRP X310, thus, drives four antennas through its two daughterboards A and B, while each USRP N210 drives two.

To make this configuration possible, we employed 64 identical 100 ft long cables to connect 32 daughterboards housed in the radio rack to 32 *antenna pairs* hung off the ceiling. The same cable lengths ensure identical electrical signal delays from the radios to the antennas and vice versa, which ultimately guarantee symbol-level synchronization across the whole testbed.

The antenna grid floor plan layout is shown in Fig. 5. The 64 antennas are deployed across 8 rails, each hosting four equidistant *antenna pairs* (8 antennas). With a spacing of 5 ft between *antenna pairs*, and 12 ft between rails, Arena covers an overall deployment area of 2240 ft$^2$ and features inter-antenna distances ranging from 5 ft to 85 ft. With respect to Arena's previous configuration presented in [7], USRPs X310 are deployed in the middle of the grid and stretch along the whole deployment length. These are more powerful SDRs with respect to USRP

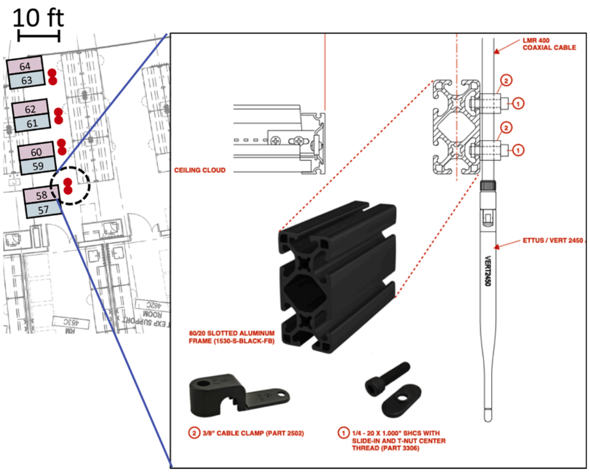


**Fig. 6.** Arena's hanging rails. Arena's 64 antennas are securely tied to 8 hanging aluminum rails via two cable clamps. Antennas can slide and be relocated along the whole length of the rail for additional configurations.

N210 and are suitable for LTE and 5G NR implementations, among others. The current configuration thus facilitates large scale cellular technologies assessment, allowing the deployment of up to 8 indoor cells along 84 ft of contiguous Line-of-Sight office space.

Moreover, being antennas hung off the ceiling, Arena offers unique Line of Sight (LOS) conditions with respect to similar deployment environment testbeds, thus guaranteeing large scale and long distance communications in a real office environment. Its grid layout eases topology changes as well as long-distance communications without requiring physical relocation of the radios. Arena scale, distances, and covered area are unprecedented for an indoor wireless deployment.

Despite of the multitude of radio testing topologies that the 64-antenna grid offers, additional topologies can easily be deployed on-demand thanks to the sliding mounts that allow each antenna to be relocated at any location across the rail, for application-specific scenarios, as illustrated in Fig. 6. The latest Arena configurations will always be available on the Arena website[1].

## 3. Hardware and software components

In this section, we provide a detailed description of Arena's design choices, as well as of its hardware and software configuration hopeful that this can be a tutorial for future wireless testbed developments.

### 3.1. Access system configuration

Arena is an open upon-grant platform, which any academic or industry researcher can access upon receiving a Northeastern University College of Engineering (COE) sponsored account. The use of Arena is regulated through the COE gateway through a Secure Shell (SSH) connection to the Arena Gateway. From the latter, it is possible to connect via SSH to the twelve Arena servers, namely `wineslab01-12`. Upon logging in, the servers automatically mount the account holder's network disk space through the NFS protocol. This guarantees users access to their files and installed programs on their network disk space on any of the 12 servers. Users have full permissions on contents located in their `home` directory, while permission to read, write, and execute other users' files is upon specific users grant, and denied by default. This guarantees users' privacy, files security, and keeps the single users' dependencies isolated.

### 3.2. Server rack configuration

The server rack is composed of 12 servers, the Arena Gateway, and a top-of-rack switch. The servers are Dell EMC PowerEdge R340 machines, namely `wineslab01-12`, running Ubuntu 16.04 LTS with `4.15.0-50-generic` Linux kernel, and provide the computational power to drive the 24 SDRs. Specifically, each server has a 6-core (12-threads each) Intel Xeon E-2186G processor with 12 MB SmartCache, 3.80 GHz base frequency (max 4.70 GHz) and four 8 GB DDR4-2666 RAM with 2666 MT/s speed. Each machine is equipped with four 8 GB DDR4-2666 RAM with 2666 MT/s speed, PERC H330 RAID controller, and a 480 GB 6 Gb/s-SATA 6 solid state drive.

Since high-speed connection and low-latency control are the keys to efficient software baseband processing, each server employs an additional Intel X520 Dual Port 10 Gigabit DA/SFP+ network card to communicate with the radios. This additional network card adds two network interfaces leveraging the SFP+ technology to establish a fast and reliable link to the radios and an aggregate data-rate of 20 Gbit/s.

A set of open-source software tools has been pre-installed on the servers to communicate with the SDRs and drive them. Among these, are GNU Radio 3.7.13, srsLTE, UHD 3.14, Python 2.7, and Python 3.5, which are ready to be used by any user to run wireless experiments. A standard Gigabit Ethernet interface connects each server to the Arena Gateway through a top-of-rack 24-Port Netgear GS324 Gigabit Ethernet switch. The gateway is implemented on a Dell Precision 5820 machine running Ubuntu 16.04 LTS with `4.15.0-50-generic` Linux kernel and implements server access control and network security features, as well as it provides Internet access to each of the 12 servers.

Powering the server rack might require a high power supply, which potentially varies over time. To this end, the power supply system is based on two APC Metered Rack Power Distribution Units (PDU) AP7811B and a Dell 5000 VA 208 V Smart Uninterruptible Power Supply (UPS). This protects all the server rack devices from power spikes and surges and provides approximately one hour of emergency power in case of outages. Moreover, the UPS is powered through an emergency power receptacle, active even in case of a power outage in the building, as a second level of power outage protection.

### 3.3. Radio rack configuration

The radio rack houses 16 National Instruments (NI) USRP N210, eight NI USRP X310, four NI OctoClock clock distributors, and two 10 Gigabit switches. USRPs are experimental hardware radio platforms completely controllable through software programs. They embed a Field-programmable Gate Array (FPGA), Analog-to-Digital Converters (ADCs), and Digital-to-Analog Converters (DACs) and are particularly suitable to design, test, prototype, and deploy wireless radio communication systems and protocols.

The USRP N210 is a networked device with high-bandwidth and dynamic range processing. It embeds a daughterboard slot allowing for bi-directional wireless communication and is controllable via software through its Gigabit Ethernet interface. Specifically, it includes a 100 MS/s Xilinx Spartan 3A-DSP 3400 FPGA, a 14-bit 100 MS/s dual ADC, a 16-bit 400 MS/s dual DAC. The host sampling rate is up to 50 MS/s, while its internal clock rate equals to 100 MHz. Furthermore, its modular design allows it to combine and synchronize multiple USRPs N210 for more advanced MIMO applications via a MIMO cable or external synchronization.

The USRP X310, instead, is a high-performance scalable device embedding two daughterboard slots as well as a user-programmable FPGA. The two daughterboard slots, allowing for two bi-directional transmit-receive chains, make it particularly suitable for MIMO applications. This device can be controlled via software through two 200 MS/s aggregate SFP+ slots. Specifically, the USRP X310 includes a 200 MS/s

[1] https://express.northeastern.edu/arena/.

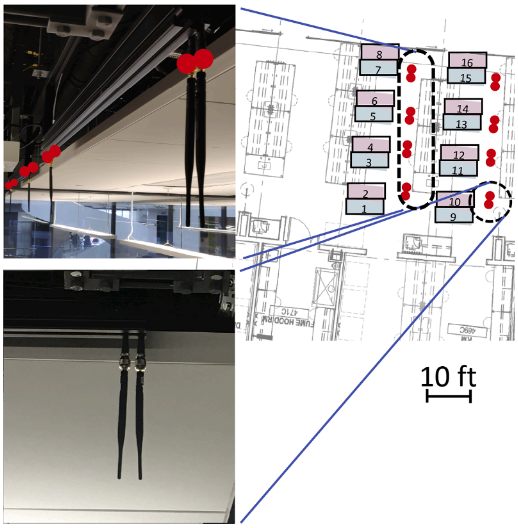


**Fig. 7.** Hanging Antennas.

XC7K410T Kintex-7 FPGA, a 14-bit 200 MS/s dual ADC, a 16-bit 800 MS/s dual DAC. The host sampling rate is up to 200 MS/s, while its internal clock rate is 200 MHz. As for the USRP N210, multiple X310 can be synchronized through external synchronization for MIMO applications. Both USRPs N210 and X310 operate from DC to 6 GHz.

Finally, USRPs N210 and X310 are fit with one and two CBX daughterboards for communication purposes, respectively. A CBX daughterboard is a full-duplex, double-chain wideband transceiver allowing USRPs to operate in the 1.2-6 GHz frequency range with up to 120 MHz instantaneous bandwidth (note that the USRP N210 can only process 56 MHz) and up to 22 dBm transmit power. The CBX daughterboard embeds one TX/RX chain, allowing the signal transmission or reception, as well as one RX2 one, allowing signal reception only. The CBX can serve a wide variety of application areas in the sub-6 GHz spectrum domain, including Wi-Fi research, LTE cellular base stations, cognitive radio research, and radar [12]. We present some of the possible research applications on this fundamentally important spectrum portion in Section 5.

### 3.4. Antenna grid configuration

One of the highlight design choices of Arena is the 64-antenna ceiling grid concerning an $8 \times 8$ array that covers an overall area of 2240 ft$^2$, where each antenna is cabled to an antenna terminal in the radio rack. Arena's Antenna Grid is based on 64 American Wire Gauge (AWG) RG8-CMP [13] low-attenuation, fireproof 100 ft long and 0.4 in thick SMA-to-SMA connection cables specifically designed for indoor applications. These are made of 2.74 mm-diameter solid annealed bare copper conductors insulated with 7.11 mm-diameter cellular fluoropolymer. They have a 2.95 Ω/km, 9.58 Ω/km shield direct current resistance, $(50 \pm 3)$ Ω impedance, 83% propagation speed, 300 V voltage rating, and 16.2 GHz cutoff frequency.

The use of 64 same-length cables guarantees equal electrical signal delays in transmission and reception across the whole testbed, independent on the location of the antenna that is transmitting or receiving. The guarantee of identical delays radios-to-antennas and antennas-to-radios is crucial for communication applications such as MIMO, beamforming, and coordinated access schemes, where different antenna connector lengths might compromise the communication effectiveness. The 100 ft cable length has been selected as the shortest distance connecting the farthest antenna to the Radio Rack so as to reduce the overall signal power loss along the cable, which equals to 6.8 dB and 11 dB over 100 ft length at 2.4 GHz and 5.0 GHz, respectively. Despite this loss is non-negligible it can be however easily compensated in software tuning the transmission and reception gains of the wireless communications. Guaranteeing consistent delays across the whole testbed comes at the price of having tens of extra feet of cumbersome cables to bundle over the ceiling of an office space. These have been professionally coiled across the ceiling avoiding loops that might result in undesirable electromagnetic effects. For over-the-air communications, Arena features 64 vertically oriented VERT2450 antennas that have been positioned upside-down to provide ground coverage as shown in Fig. 7. This unusual deployment changes the antennas polarization. While for remote users this would not impact their performance evaluation, the on-site users are invited to consider this aspect when employing additional radio devices for experiments. antenna deployment inverts the antenna nau The VERT2450 is a 3 dBi-gain toroidal-radiation dipole dual-band antenna, optimized to work in the $2.4 - 2.5$ GHz and $4.9 - 5.9$ GHz frequency bands and offer a 50 Ω nominal impedance [14]. The VERT2450 antenna's patterns are reported in Fig. 8.

The 64 VERT2450 antennas are mounted through eight 15 ft rails hanging off the ceiling. Rails have a $1.5 \times 3$ inches rectangular T-slotted profile with six open slots on each of its sides and one on the front side.

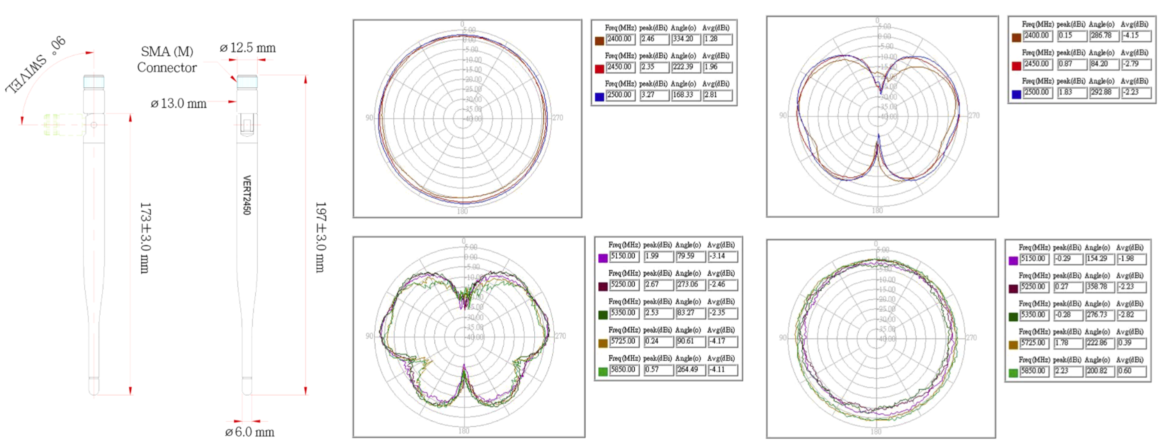


**Fig. 8.** Arena's antennas and radiation patterns. Arena features omnidirectional dual-band VERT 2450 antennas pointed downwards from the ceiling. We herein illustrate the Horizontal and Elevation plane antenna patterns at $2.4 - 2.5$ GHz and $5.150 - 5.850$ GHz bands.

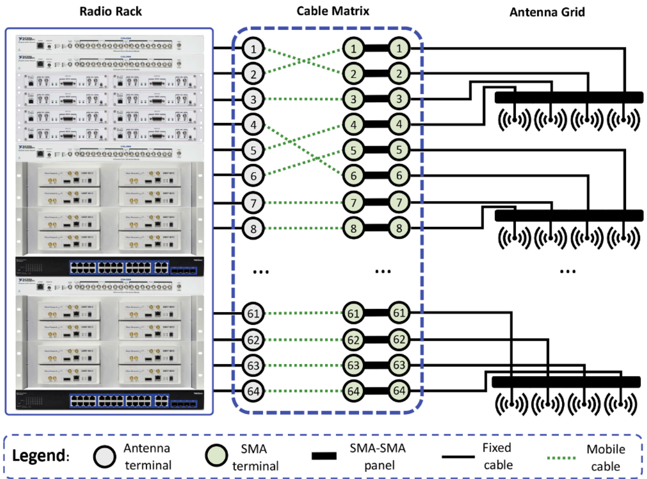


**Fig. 9.** Arena's Cable Matrix.

Each rail hosts four *antenna pairs* spaced 5 ft from each other, while the same pair antennas are spaced 1 inch only so as to reproduce the same antenna spacing at the output of an USRP X310 and USRP N210 SDR enclosure. The two antennas in a pair are connected to the TX/RX and the RX2 chains of a single daughterboard, respectively. The two antennas can, in principle, both be used as receivers (even though the physical proximity results in limited spatial diversity in the sub-6GHz regime) but not both as transmitters. Their recommended configuration, thus, is to operate in full-duplex mode where the antenna connected to the TX/RX daughterboard terminal is employed to transmit while the one connected to the RX2 terminal is used to receive, e.g., to mimic a full-duplex cellular user. Last, it is worth mentioning that the limited spacing between the two antennas might result in undesired shadowing effects when receiving signals from other antennas along the rail. This effect is however negligible when involving antennas belonging to different rails. The users are thus encouraged to avoid same-rail radio communications to avoid this effect. The rail's robust structure provides enough strength to support the weight of eight AWG cables each, while their modular design permits antennas relocation along their whole length. Indeed, antenna locations can be effortlessly adjusted by just sliding them along rails as previously mentioned and shown in Fig. 6.

### 3.4.1. *The cable matrix*

Compared to the design presented in [7], we introduce the possibility of switching the one-to-one mapping between Radio Rack antenna terminals and Antenna Grid antenna points. To make this possible, we introduce Arena's Cable Matrix. This is a 64-SMA-input 64-SMA-output matrix placed in between the Radio Rack and the Antenna Grid as shown in Fig. 9. The Cable Matrix consists of a set of two-sided panels, featuring 64 low-loss SMA connectors each side and 64 5 ft long identical cables. We employed the same American Wire Gauge (AWG) RG8-CMP low-attenuation, fireproof cables as the ones used for the Antenna Grid [13]. Once again, same-length cables guarantee identical delays at all the radios, which is fundamental for full-testbed synchronization.

The front of the Cable Matrix is connected to the Antenna Grid through the 100 ft cables presented in Section 3.4, while the back connects to the Radio Rack through its 64 5 ft-long cables. The function of this simple, yet effective setup is to make it possible to virtually switch the locations of transmitters and receivers, thus enabling swift topology reconfiguration. This feature enables a variety of use-cases and applications that require the very same device boards to transmit or receive from multiple physical locations. An extensive experimental assessment of these technologies would require physical radio relocation. Instead, Arena's Cable Matrix makes it possible to change the location of transmitting and receiving radio boards selecting among 64 different choices, by simply switching the 5 ft cables interfacing the Radio Rack. We present a series applications that employ the Cable Matrix as a fundamental part of their assessment in Section 5.3.

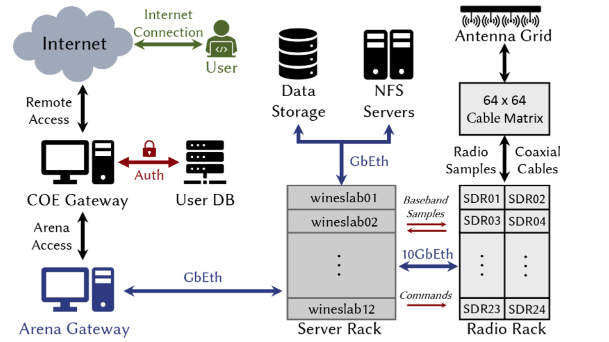


**Fig. 10.** Arena access system diagram.

## 4. Life-cycle of an experiment

To access Arena and perform real-time experiments, external users can request a Northeastern University College of Engineering (COE) sponsored account to the WiNES Lab mailing list owner[2]. Upon getting a sponsored account, users can authenticate to the WiNES Lab Gateway[3] through the COE Gateway[4] which is accessible from the Internet. Upon accessing the COE Gateway, granted users will be allocated dedicated network disk space, accessible through any machine under the COE domain via the Network File System (NFS) protocol. Once logged into the WiNES Lab Gateway, it is possible to SSH to any of the 12 Arena servers[5] described above and create files, develop code, install desired software, and run real-time experiments with the guarantee of a secured file system. We report a diagram of the access system in Fig. 10.

We have seen in Section 2.2 how each `wineslab01–12` server drives only one or a subset of the available SDRs, therefore, the server selection has to be considerate of the user's experiment of choice. Users can, thus, access one, some, or all of the 12 servers at once, depending on their needs. Referring to Fig. 4, servers `wineslab01–08` drive one and only one USRP X310 per server, namely SDR 1–SDR 8, while servers `wineslab09–12` drive four USRPs N210 each through their two 10 Gigabit network interfaces. Specifically, servers `wineslab09`, `wineslab10`, `wineslab11`, and `wineslab12` drive SDR 9–SDR 12, SDR 13–SDR 16, SDR 17–SDR 20, and SDR 21–SDR 24, respectively. Finally, Arena Gateway keeps track of logged users to avoid server overloading and possible conflicts of users driving the same radios.

Despite users being able to install their software of choice on the network disk space allocated to them, basic development and testing software, such as Python, GNU Radio, and srsLTE, has already been installed on the servers and it is accessible by all users. Performing real-time wireless experiments is as simple as running pre-compiled software like GNU Radio. GNU Radio features sample physical layer pre-compiled programs such as `benchmark_tx.py` and `benchmark_rx.py` that can be run to implement different transmission techniques such as narrow-band and OFDM with a long list of tunable physical-layer parameters such as operational bandwidth, transmission and reception gain, and modulation. GNU Radio can implement single point-to-point wireless

[2] wineslab-owner@coe.neu.edu.
[3] wineslabgate.coe.neu.edu.
[4] gateway.coe.neu.edu.
[5] wineslab01-12.coe.neu.edu.

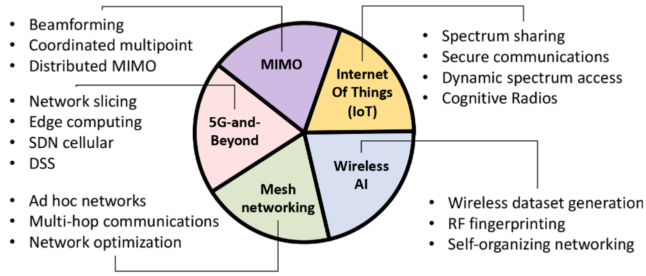


**Fig. 11.** Wireless technologies applicable on Arena.

links or more complex networking stacks, including multi-hop connectivity, wireless mesh networks, and MIMO communications. More advanced pre-compiled software can also seamlessly run on any Arena servers. For example, srsLTE is an open-software standard-compliant full stack implementation of LTE and 5G NR. Any user can easily instantiate a softwarized LTE/NR core network and LTE/NR base station by simply running the `srsepc` and `srsenb` programs, as well as instantiate LTE/NR users through the command `srsue`. GNU Radio and srsLTE instances can be run at any given antenna locations by accessing the desired Arena server and addressing the intended Arena SDRs. This can be done by GNU Radio command-line tools (e.g., "–args="addr=$SDR-9"" to drive SDR 9) or by editing the srsLTE configuration files accordingly. Explicit radio addressing is not needed for USRP X310, which are one-to-one driven by servers `wineslab01–08`, however users can specify the intended daughterboard to use, e.g., via "–args="addr=$SDR1" –spec "A:0"" to drive SDR 1 daughterboard A.

The USRP Hardware Driver (UHD), a user-space library that runs on a General Purpose Processor (GPP), handles the baseband samples between the SDR and the control host. The control host software (e.g., GNU Radio, srsLTE), will report network status and measured metrics to the user such as transmitted, received, and correctly decoded packets, network throughput, and nodes interference levels, which can be saved on network disk or external drive for further analysis. To terminate an experiment session, the user can simply log out from the servers in use and from the Arena Gateway.

## 5. Experimental capabilities

Throughout the last decade, the telecommunication industry witnessed a clear paradigm shift toward the "softwarization" of traditionally hardware-implemented functionalities. This trend promoted the strong candidacy of Software-Defined Radios as an alternative to closed-hardware vendor-proprietary solutions for wireless communications. This, in turn, produced a tremendous increase in the availability of open-source and license-free software implementation of standard protocols for SDRs, both from international research bodies and private-public consortia. As of today, all the major standard wireless technologies can count on a publicly released license-free SDR implementation. Among these, it is worth mentioning IEEE 802.11 (Wi-Fi) [15], LTE-A [16,17], 5G NR [18], IEEE 802.15.4 (Bluetooth) [19], LoRa [20,21], IEEE 802.11p (wireless access in vehicular environments) [22], and GNSS [23–25]. In this sense, the number of different wireless technologies that are implementable on Arena is potentially unlimited. A selection of relevant technologies is shown in Fig. 11.

Ultimately, Arena's real-time testing capabilities can be employed to experiment with point-to-point links, test multi-hop transmissions, evaluate 5G-and-beyond cellular network performance, or implement MIMO communication schemes, both for real-time research validation and live demonstrations. In the following, we extensively showcase a series of timely wireless technologies implementation and evaluation on Arena, which includes published work that employed Arena for its experimental assessment, as well as other wireless developments.

### 5.1. MIMO Communications

Among Arena highlights is its full-testbed symbol-level synchronization. This can be employed to implement synchronization-based communication techniques such as MIMO, cooperative multi-point, beamforming spectrum access, and distributed MISO, among others. Among the works using such techniques is CoBeam [26], which proposes to solve the challenges of spectrum sharing in unlicensed bands via *cognitive beamforming*. Specifically, CoBeam investigates the spectrum access problem in unlicensed spectrum bands populated by a primary and a secondary technology, where the secondary users opportunistically access the shared spectrum bands yet making sure to not interfere with the primary users' activity. CoBeam proposes a novel technique branded *cognitive beamforming* where a multi-antenna equipped secondary user accesses the shared spectrum by precoding its transmission so as to: (i) deliver the intended information to the intended receiver, (ii) null the interference at one or multiple primary users present on the channel. This way, the activity of the secondary technologies does not interfere with the primary system, making it possible for both primary and secondary technologies to access the spectrum at the very same time. Further, CoBeam implements different beamforming scheme choices so as to adapt to the traffic-varying channel conditions, including Zero-Forcing beamforming (ZF) and Maximum Ratio Transmission (MRT).

CoBeam operates at a multi-antenna secondary transmitter, and requires tight synchronization among its all transmission chains. To assess CoBeam's performances, the authors employed Arena's full-testbed symbol-level synchronization. Specifically, the Authors prototyped CoBeam on a 4-antennas secondary LTE transmitter installed at SDR 1A-1B-3-4[6] They assessed the network performance of CoBeam by prototyping a co-located primary Wi-Fi network composed of a Wi-Fi Access Point implemented at SDR 3A and a Wi-Fi station implemented at SDR 10, while the secondary transmitter attempts to deliver data to a single-antenna LTE receiver.

The Authors in [26] assessed CoBeam at a 4-antenna device comparing it with a single antenna transmitter. They switched the location of the secondary U-LTE receiver for a total of 18 different network topologies, namely employing SDR 18-2A-2B-19-3B-20-4A-4B-12 as shown in Fig. 12a. On Arena, changing the topology configuration of network deployment is fully software-controllable, and does not require any radio relocation. Average aggregate network throughput for over ten 1-minute long experiments are presented in Fig. 12b. For all the considered topologies, CoBeam achieves higher aggregate network performance than the single antenna transmitter, with average aggregate network throughput improvement of 169%.

### 5.2. 5G-And-beyond cellular networks

In the last years, research bodies and industry partners such as telco operators and radio vendors merged their efforts toward the development of open-source code for cellular networks. These efforts produced full-stack-compliant implementations of standards such as 3GPP 5G NR for general-purpose hardware like SDR [16,17]. These open-source, available-to-all developments offered a series of unprecedented research opportunities to academics and scientists around the world. At the same time, they helped telco operators to shift towards the softwarization of cellular networks, paving the way for a series of cross-industry, collaborative research alliances, such as the Open Network Automation Platform (ONAP), and O-RAN [27,28]. As a result,

[6] USRP X310 daughterboard A and B drive different antennas as shown in 5. E.g., SDR 1A drives antennas 3–4 while SDR 1B drives antennas 5–6. The same holds for the other USRPs X310.

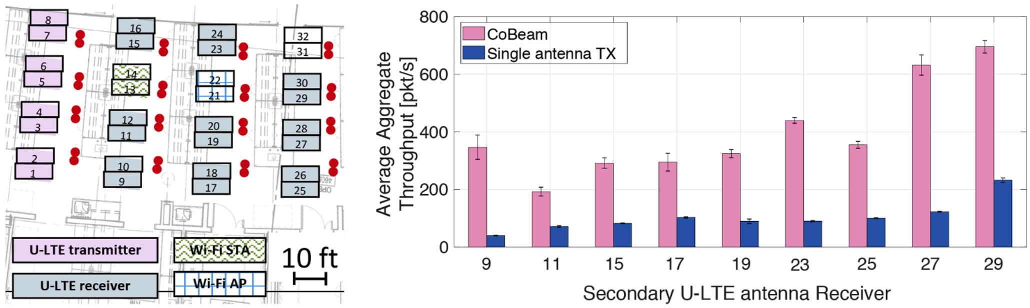


(a) CoBeam evaluation topology on Arena. Different U-LTE receiver locations are used throughout the evaluation.

(b) CoBeam aggregate network throughput performance when employing Zero Forcing (ZF) beamforming against single antenna spectrum access.

**Fig. 12.** CoBeam assessment on Arena.

not only telco operators and vendors, but anybody in control of an SDR as cheap as \$900 can instantiate a full-stack standard-compliant 5G NR gNB, and test their technological developments. To connect users to these "softwarized" cells, developers can use commercial off-the-shelf cellphones or SDR-based implementations of Users Equipment (UEs). This way researchers have an unprecedented tool to develop and assess novel cellular network developments such as dynamic spectrum access, optimal interference mitigation, network slicing, and edge computing, among others. With its scale and its system accessible all around the world, Arena represents a one-of-a-kind testing platform to experimentally evaluate newer generation indoor cell developments in a realistic wireless environment. In this section, we present some key 5g-and--beyond technologies implementation and evaluation on Arena.

*5.2.1. Multi-cell 5G cellular networks*

We start showcasing a multiple-cell 5G network on Arena, where each gNB serves multiple users. Instantiating one or more gNBs on Arena is as simple as accessing the intended server(s) and running pre-compiled software such as srsLTE [16]. As a proof of concept, we instantiate two gNBs on Arena, while we use commercial off-the-shelf (COTS) cellphones as users. In our experiments, each gNB serves three cellular users, implemented through Samsung Galaxy S5 cellphones, which request downlink traffic at full buffer capacity. The gNBs operate in a Frequency Division Duplex (FDD) configuration on NR band 7. We overall consider two network deployments scenarios as illustrated in Fig. 13: (i) A high-interference scenario where the two gNBs are relatively close to each other, and (ii) a low-interference scenario where the two gNBs are more spread out.

In the former, the gNBs are near and have largely-overlapping coverage areas. This case aims to study the interference consequences of dense deployments, such as concert venues or fairs, where mobile operators co-locate temporary pico- and femto-cells with the fixed ground infrastructure, so as to sustain the temporary increased traffic demand. On the other hand, the latter sees gNBs more spread out with only partially overlapping coverage areas, resulting in only slight inter-cell interference. On Arena, we instantiated gNBs on SDRs 7A-8A for the high interference scenario, and SDRs 5B-8B for the low interference one, as shown in Fig. 13a and Fig. 13b, respectively.

We measured the performance of the 2-gNB network in terms of average user experienced service and overall network throughput for the two deployment scenarios, as reported in Fig. 13c. As expected, lower inter-cell interference results in better users' service, with throughput gains as high as 22.84 Mbit/s (15.20 Mbit/s on average).

*5.2.2. 5G-Slicing*

In the context of network virtualization and softwarization, network slicing will be the cornerstone of 5G networks and the Internet of Things. This technology allows multiple Mobile Virtual Network Operators (MVNOs) to share a common underlying physical infrastructure and dynamically deploy "slices" tailored for specific services (e.g., video streaming, augmented reality) or requirements (e.g., low latency, high throughput, low jitter) [29,30]. By leveraging network sharing and on-demand instantiation of a virtualized network infrastructure, network slicing avoids static—and frequently inefficient—network deployments that have plagued traditional hardware-based cellular networks in the past.

Among the most timely applications, is the cellular Radio Access Network (RAN) slicing. Through RAN slicing, the cellular infrastructure

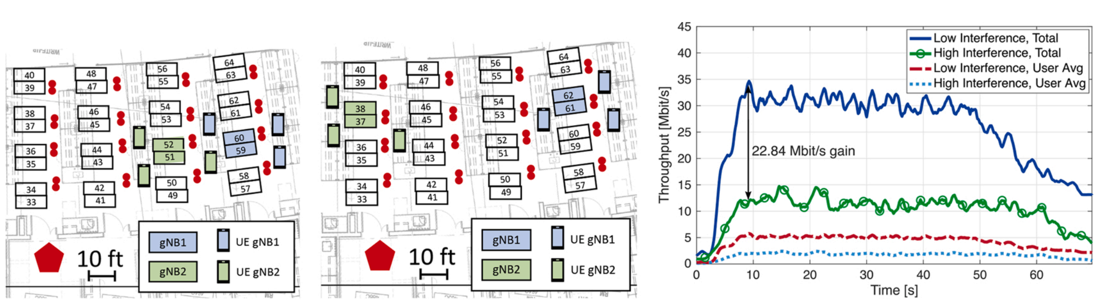


(a) High-interference cellular evaluation topology.

(b) Low-interference cellular evaluation topology.

(c) Network throughput in high and low inter-cell interference scenarios using 2 gNBs and 6 UEs with srsLTE.

**Fig. 13.** Cellular network assessment on Arena.

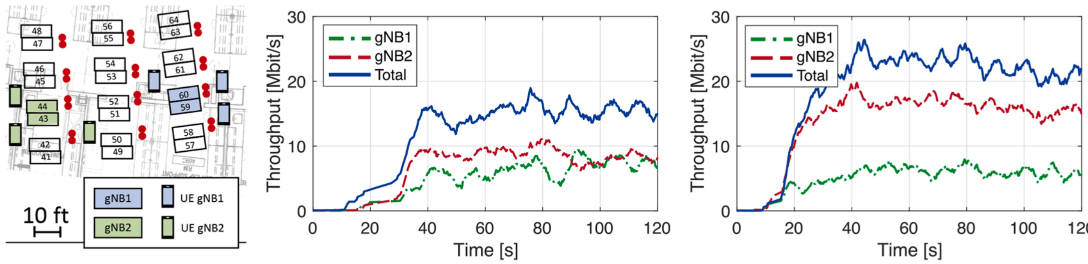


(a) Cellular slicing evaluation topology. (b) Cellular slicing w/o isolation among slices. (c) Cellular slicing w/ isolation among slices.

**Fig. 14.** Cellular network slicing assessment on Arena.

resources are allocated to multiple MVNOs that dynamically instantiate virtual cellular networks [31]. Among its benefits, RAN slicing allows MVNOs to deploy cellular networks on-the-fly, to supply to time-varying traffic demand, user location, and Quality-of-Service (QoS) requirements.

We here showcase a RAN slicing implementation on Arena. We study the scenario where two MVNO lease spectrum resources from an infrastructure provider, so as to instantiate RAN slices and serve their cellular users. In our example, we assume the two MVNOs requesting the same amount of resources and we analyze the performance of two network slicing policies adopted by the infrastructure provider. Specifically, we consider the policy presented in [31], where spectrum resources are allocated to reduce the cross-slice interference (i.e., slice isolation), and pseudo-random interference-agnostic allocation policy. On Arena, we leverage srsLTE to instantiate two gNBs, at SDR 8A and 6A, as shown in Fig. 14a. Each gNB operates in FDD mode on NR band 7, and provides service to six cellular users, implemented with COTS Samsung Galaxy S5 smartphones. We define an overall useful bandwidth of 10 MHz, corresponding to 50 Physical Resource Blocks (PRBs) while each RAN slice corresponds to 25 PRBs (5 MHz). We assess the performance of the two policies presented above in terms of average per-user throughput and overall network throughput. The network performances for a single run of our experiments are reported in Fig. 14b and Fig. 14c, respectively. On average, RAN slice isolation improves the per-user experienced service by 30% (approximately 6 Mbit/s) with gains up to 9 Mbit/s.

### *5.2.3. Mobile edge computing*

By combining computational, networking, and storage capabilities, Arena is the perfect candidate to prototype and evaluate Multi-access Edge Computing (MEC) solutions. MEC brings network functionalities and services to the edge of the network, thus enabling technologies like edge caching and content delivery, geographical video streaming, and agile instantiation of small cells to support, for example, Augmented and Virtual Reality (AR/VR). Among other benefits, moving intelligence in the proximity of the users reduces communication latency and facilitates the implementation of high-speed cellular communications. For these and other reasons, MEC has been identified as one of the pillars of 5G systems.

One of the MEC solutions for 5G systems is Sl-EDGE [32]. This is a framework for MEC ecosystems that enables simultaneous instantiation of networking, computational, and storage services. In Sl-EDGE, network resources are abstracted and represented as virtualized edge nodes. This way, edge nodes constitute a shared infrastructure among service providers, which can employ, for example, RF resources to provide connectivity to mobile users (e.g., Wi-Fi, LTE) and CPU resources to implement services like caching, video trans-coding, and in-network computing.

To demonstrate the applicability of multi-access edge computing to real wireless systems, we have prototyped Sl-EDGE on Arena. Our experiments aim to demonstrate how intelligent edge resource allocation can abstract the wireless and computational resources of a distributed infrastructure, and successfully instantiate diversified on-demand services across the network. Our prototype consists of 14 edge nodes, each one including one radio and one server. Each server implements one (or more) computational/storage services as well as drives performs the baseband processing of one (or more) radios supporting their wireless communication over-the-air. Specifically, edge servers implement *dash.js* video streaming, *ffmpeg* video trans-coding, and caching services, while they drive radios implementing LTE base stations running srsLTE (SDR 2B-4B-6B-8A-16-18-20-22-24) and Wi-Fi terminals through GNU Radio (SDR 15-17-19-21-23). Lastly, LTE base stations are intended to serve cellular users, for which we used COTS smartphones, while Wi-Fi terminals communicate with each other in an ad-hoc manner. The ultimate goal of Sl-EDGE is to allocate heterogeneous network services at the edge nodes with the guarantee of respecting the single-node radio and computational/storage constraints.

Figures. 15 b and 15 c report Sl-EDGE resource allocation across the

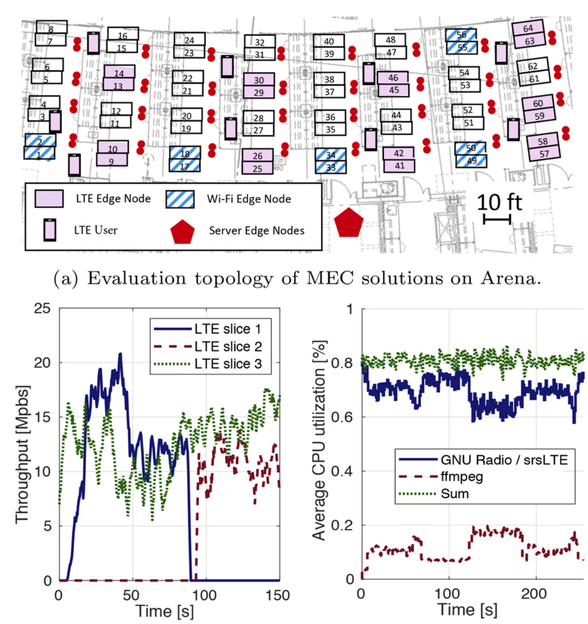


(a) Evaluation topology of MEC solutions on Arena.

(b) Per-slice throughput experienced by LTE users with dynamic Sl-EDGE solution. (c) Edge servers CPU utilization. Sl-EDGE respects single-node constraints.

**Fig. 15.** Sl-EDGE performance assessment on Arena.

network when 11 heterogeneous slices run on Arena. Specifically, Fig. 15b illustrates the dynamic allocation of LTE slices (slices 1–3) on the same edge radios. At the same time, the edge servers implement video transcoding (slices 4–8) and video streaming services (slices 9–11). The curves in the figure report the per-slice instantaneous cellular user experienced service over time. Fig. 15c, instead, reports the edge servers' resource utilization. While allocating heterogeneous computational services on the shared network infrastructure, Sl-EDGE guarantees their correct implementation, as shown by the servers' CPU utilization over time in Fig. 15c.

### 5.3. Wireless artificial intelligence and real-Wireless dataset generation

As 5G and the IoT will require unprecedented data rate levels in the sub-6 frequency bands, the design of fine-grained spectrum analysis and optimization techniques has become a timely and urgent necessity [33, 34]. To this end, the design and development of techniques able to extract knowledge from the spectrum in real-time and select the optimal spectrum access strategy accordingly has become more important than ever. Recently, *wireless deep learning* has been proven to be extremely successful in addressing problems such as modulation recognition [35], radio fingerprinting [36] and medium access control [37], and has taken us many steps in the right direction [38,39]. Thanks to its unique theoretical and practical advantages, deep learning has been exceptionally successful in addressing classification and optimization problems where closed-form mathematical expressions are difficult or impossible to obtain [40]. As artificial intelligence wireless application grows in popularity, the need for publicly available real-wireless data-set is more stringent. In the following, we report a series of experimental AI wireless applications using Arena.

#### 5.3.1. RF-Fingerprinting

Radio fingerprinting has lately received significant attention as a reliable identification technique for wireless devices[36]. This technique leverages the fabrication hardware imperfections of a transmitter's radio circuitry to discriminate among nominally identical transmitting boards. These imperfections can be exploited at the receiver at the physical layer by analyzing the received wireless signals and thus authenticate the transmitter. This way, the authentication can happen right at the waveform layer with no coordination or message exchange – thus avoiding energy-hungry cryptography techniques. This concept is further exploited by applying deep learning algorithms, which has been demonstrated to significantly enhance the classification accuracy of radio fingerprinting.

The implementation of a fully functional radio fingerprinting system, however, faces many challenges typical of wireless systems. Perhaps one of the most crucial challenges is understanding how the non-stationary wireless channel impacts the fingerprinting accuracy. Through over-the-air communications, indeed, the channel action is superimposed to the wireless signal, thus "blurring" the transmitter's hardware impairments at the receiver and undermining device identification effectiveness. The authors in [41] aimed to dissect the impact of wireless channels on radio fingerprinting by studying the relationship between Wi-Fi radio signal classification and the relative positioning of the transmitter (the device to be classified) and the receiver (the classifier). Specifically, through an extensive experimental campaign, they investigated whether the cognitive neural network (CNN) employed at the receiver's classifier was invariant of the non-stationary wireless channel. In doing so, [41] provides a systematic and quantitative evaluation of the impact of the wireless channel on CNN-based radio fingerprinting algorithms performance.

The authors in [41] employed Arena to instantiate 20 802.11a/g (WiFi) transmitters, and one WiFi receiver intended to classify the transmitter's signals, as shown in Fig 16 a. The authors overall collected 5 TB of raw and equalized Wi-Fi IQ-samples[7] for every single transmitting board over 10 different days, and thus wireless channels, for a total of 240 unique transmitter board-day pairs.

Then, they performed CNN training and classification offline, mixing and matching training days with testing days, so as to assess the performance of the classification at the variation of the channel characteristics. The classification results for 20 transmitters over 10 different days are reported in the confusion matrix in Fig. 16b. In Fig. 16b, the matrix's axes represent the testing and the training days, respectively, while the matrix diagonal entries are same-day training and testing. The performance results highlight the non-independence relationship between radio fingerprinting accuracy and wireless channel, scoring an overall classification accuracy of 82% when train and test happen on the same day and 5% only when they happen on different days.

Moreover, in [41] the authors employed Arena to switch the antenna-radio mappings as described in Section 3.4.1. To do so, they used the possibility of employing one designated antenna for all the transmissions such that all transmitter boards were equally distant from the receiver and experienced similar channel conditions as shown in Fig. 16c. As for the previous experiment, they collected raw IQ-samples for every transmitter-receiver pair and they assessed the classification accuracy of the CNN at the variation of the transmitter board number only. When testing the CNN's performance for different transmitters with the same channel, the overall classification accuracy is 83.5% suggesting that filtering out the channel eases the fingerprinting effectiveness. As a last note, the authors will release the collected data-set to the wireless community so as to foster the research on radio fingerprinting, suggesting orthogonal uses of Arena as a wireless channel data-set generation platform.

#### 5.3.2. Lora fingerprinting

The number of connected IoT devices has increased exponentially in recent years, and it is expected to expand at a very fast pace. This dramatic growth creates the need for power-efficient, low-cost, and long-range reliable communication technologies [42]. Among other Low Power Wide Area Network (LPWAN) technologies, LoRa is one of the main candidates for this task [43]. LoRa is a long-range, low data-rate wireless protocol operating in unlicensed ISM bands, specifically designed for small, power-constrained IoT devices. As mentioned in Section 5.3.1, waveform-layer authentication techniques for wireless devices are an effective and power-efficient alternative to sophisticated security protocols. The concept is exacerbated when we focus on a large number of battery-constrained, computational-constrained IoT devices. In this section we investigate the challenges of applying RF-Fingerprinting to LoRa devices.

Outside of the wireless-related challenges presented in Section 5.3.1, RF fingerprinting of LoRa devices is most importantly undermined by their spotty and unpredictable transmission patterns. This aspect increases the complexity of the features that are exposed to the classification algorithm at every transmission, and how to tackle this problem still represents an open research topic. Toward the development of effective RF fingerprinting techniques for LPWAN devices, the data-set collection is a first fundamental step. Along this line of research, Arena can play a fundamental role in collecting large real-wireless measurements. These measurements can be then used to train and assess the performance of several deep-learning algorithms.

As a first step, we performed an extensive measurement campaign on Arena, where we employ dozens of nominally identical LoRa chip-sets and we deploy them all together in a single location as shown in Fig. 17a. LoRa devices operate at the 902.3 MHz band and transmit 100

[7] IQ-samples are bare physical layer samples that do not contain any transmitter-specific information in plain sight.

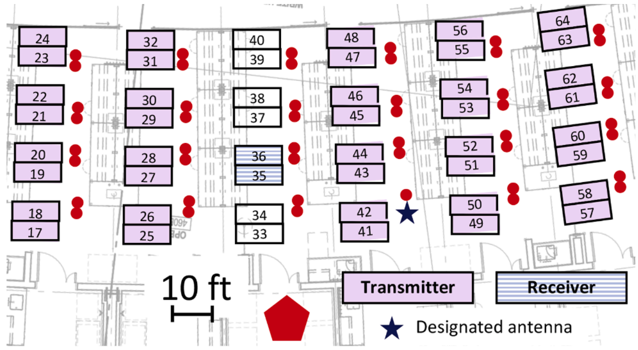


(a) RF-Fingerprinting assessment topology on Arena. 20 nominally identical radios boards transmit the exact same radio signal while the receiver attempts at identify them based on their hardware imperfections only. Radio boards are then shuffled around (simply swapping cables on Arena's matrix) to prove RF-Fingerprinting independence on the wireless channel.

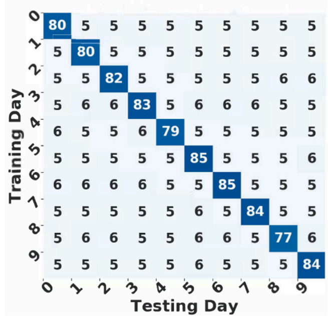


(b) CNN classifier assessment on Arena at the variation of the classifier training and testing day

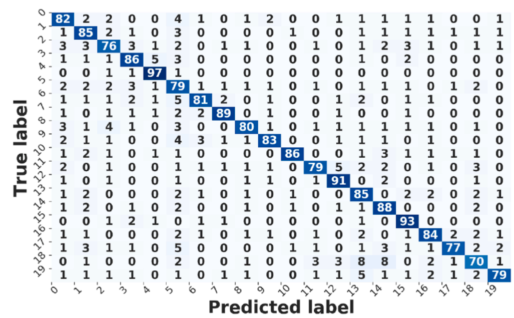


(c) CNN classifier assessment on Arena using a designated antenna for all the transmitters

**Fig. 16.** RF-Fingerprinting assessment on Arena.

12 B-packets at a time, carrying temperature, humidity, and voltage information, in a bursty fashion. In the meantime, we employ Arena radios running a GNU Radio receiver to collect raw waveform-layer measurements at several locations in space. Last, we collect measurements for different LoRa transmission parameters, such as output transmission power and spreading factor. One more time, Arena offers unique channel diversity conditions, ideal for real-wireless data-set collection. As a preliminary evaluation step of our analysis, we implemented the RF-fingerprinting algorithm presented in Section 5.3.1 and assess its performance for five LoRa devices. The classification accuracy results are reported in Fig. 17b. While the porting of the RF-Fingerprinting algorithm developed for Wi-Fi shows promising results on LPWAN (overall classification accuracy of 90%), we keep investigating other research directions for RF-based LoRa-specific device identification.

### 5.3.3. *Wireless adversarial deep learning*

We have seen in the previous sections how deep neural networks (DNN) can implement low-cost authentication techniques for wireless devices. Like other technologies, however, DNNs are vulnerable to adversarial attacks [36,44]. The possible adversarial attacks are many. One attacker, for example, can impersonate a legitimate device by spoofing its waveform to trick the classifier. This way an adversarial device can authenticate as a legitimate device and break into a protected system. Another way an attacker can pose security threats to an RF classifier is by super-imposing non-random wireless perturbations to a legitimate device waveform that is trying to authenticate. This way, the attacker can perform a denial of service (DoS) attack by "blurring" the hardware impairments of a legitimate wireless signal, which will result in a device authentication failure. Ultimately, while RF fingerprinting is a promising technique aiming to revolutionize the security of wireless devices, a robust and secure implementation of this technique has to be considerate of several security challenges.

In here, we investigate security pitfalls of RF fingerprinting and assess the robustness of this classification method to adversarial attacks. Specifically, we assess the classification performance of RF fingerprinting under two specific attacks.

(i) Adversarial Waveform, where a set of adversary devices tries to spoof their transmitted signal to imitate legitimate devices and fool the classifier, and (ii), Replay Waveform, where a set of adversary devices employs an adversary receiver to eavesdrop legitimate transmitter signals, and then imitate the legitimate devices employing this extra information to trick the receiver. This second method is more sophisticated as the sniffed signal is used to spoof the hardware impairment information of the legitimate devices.

We employed Arena flexible setup to carry out these wireless adversarial attacks and performance assessment of our fingerprinting technique. The experimental setup consists of five legitimate transmitters (in blue), one legitimate receiver implementing a CNN-based authenticator (in yellow), and five adversary transmitters (in pink), as illustrated in Fig. 18a. For this configuration, we employed different SDRs and Arena's capability to reconfigure its radio-antenna mappings as described in Section 3.4.1. All the transmitters employ an IEEE 802.11 GNU Radio implementation and act as Wi-Fi transmitters, while the receiver implements a physical-layer receiver collecting waveform-layer samples. The classification performance of the system in the absence of adversarial attacks is reported in Fig. 18b (left) and scores an average classification accuracy is 59%.

The first adversarial attack we analyze consists of five attackers tentatively impersonating one legitimate device each by transmitting

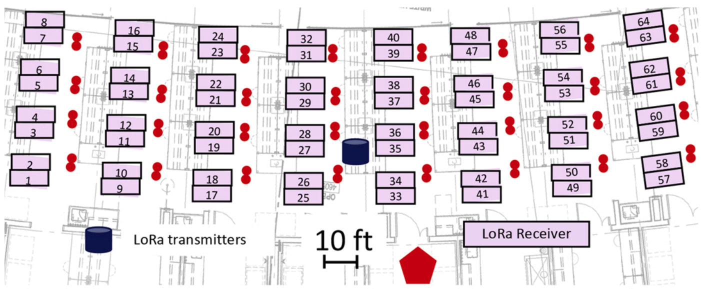


(a) LoRa-Fingerprinting assessment topology on Arena. Several of nominally identical LoRa devices are deployed in one location while Arena radios are employed as LoRa receivers.

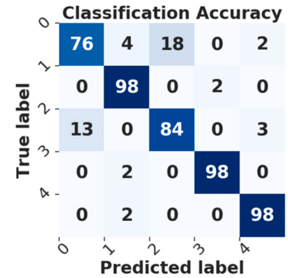


(b) Preliminary RF fingerprinting performance assessment using five LoRa devices.

**Fig. 17.** LoRa technology Fingerprinting assessment on Arena.

legitimate devices' data over-the-air. In this attack scenario, the adversary transmitters are co-located with their targeted legitimate devices so as to benefit from similar channel characteristics. In this first configuration, however, the adversary devices do not imitate the legitimate devices' hardware imperfections but just transmit the same data. The classification accuracy of the system under this attack is shown in Fig. 18b (center). Being the deployed classifier looking for specific hardware features in the wireless signal, transmitting the same data does not suffice in imitating a legitimate device. The results, indeed, indicate a weak chance impersonation, with a fooling rate of only 20%.

The second threat we pose to our system is more sophisticated. We employ an adversary receiver (in green in Fig. 18a) co-located with the legitimate receiver, so as to eavesdrop legitimate device wave-forms and pass them to adversary transmitters to trick the classifier. This way, an adversary transmitter can benefit not only of the same channel as the legitimate device but also of its hardware impairments. The performance of our system for this second attack scenario are reported in Fig. 18b (right). Despite the information gathered by the adversary receiver is a mix of hardware impairments and wireless channel, this is enough to increase the impersonation rate by 50% when compared to our previous attack.

This set of preliminary experiments on Arena helped us get a grip on the robustness of artificial intelligence identification methods for wireless devices, inspiring future research directions for AI-powered wireless Internet of Things.

### 5.4. IEEE 802.11 (Wi-Fi)

As the development of wireless communications continues, IEEE 802.11 protocol family remains the reference standard for affordable, unlicensed spectrum access. Moreover, the sub-6 GHz ISM bands used by Wi-Fi are often the default option for new technology deployment, exacerbating the spectrum coexistence issue in what is already the most crowded portion of the spectrum. As Wi-Fi chip-sets become cheaper and more accessible than ever before, the experimental assessment of this standard's developments remains of paramount importance. Thanks to their hardware flexibility and the availability of open-source implementations of IEEE 802.11, SDR is the platform of choice for Wi-Fi developments. Arena's scale and deployment environment, on this end, represent a unique testing platform for IEEE 802.11 and other technologies on the same spectrum bands.

#### 5.4.1. Cognitive-Radio (CR) wi-Fi

Dynamic Spectrum Sharing (DSS) networking solutions like cognitive radios (CR) can also be implemented on Arena. Concepts like spectrum sharing, opportunistic spectrum access, and spectrum management typically rely on information gathered by idle listening on the wireless channel [45–47]. We herein showcase a cognitive radio (CR)

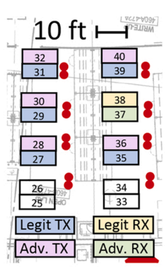


(a) Adversarial attack deployment on Arena. 5 legitimate transmitters authentication is threaten by 5 adversarial transmitter devices in scenario (i). In scenario (ii), the adversarial transmitters are helped by an adversarial receiver.

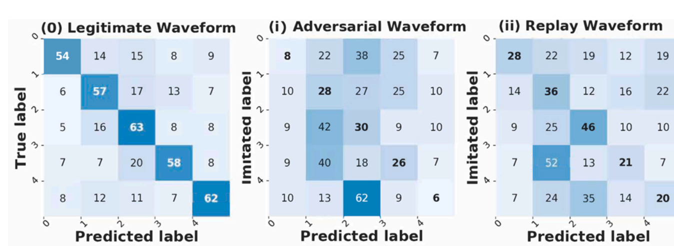


(b) Performance Confusion Matrices for different deployments on Arena. (left) Performance for the deployed system in the absence of adversarial attacks. The average classification accuracy is 59%. (middle) Performance for 5 adversary transmitters imitating 5 legitimate devices. The attacker does not succeed in fooling the classifier as the classification is based on board-specific hardware impairments. The average misidentification performance is only 20%. (right) Performance for 5 adversary transmitters imitating 5 legitimate devices with the help of an adversary receiver. The adversary receiver informs the attackers of specific-board hardware impairments of the legitimate devices raising the average misidentification accuracy to 30%.

**Fig. 18.** Wireless adversarial deep learning assessment on Arena.

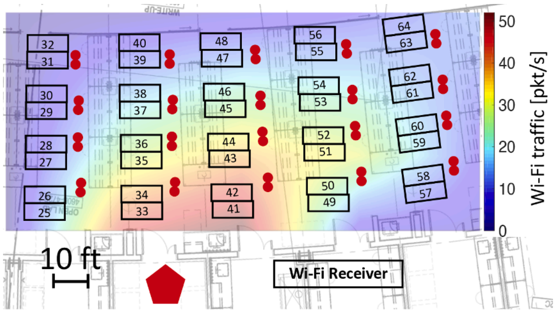


**Fig. 19.** Heatmap of sensed Wi-Fi activity at different locations on Arena.

application on the crowded 2.4 GHz ISM band, aimed at gathering information about the Wi-Fi activity in the surrounding.

As mentioned earlier, Arena is deployed in a real office environment, featuring rich channel characteristics, surrounding objects mobility, and wireless activity typical of indoor spaces, including Wi-Fi. In this implementation, we leverage Arena to passively eavesdrop the commercial Wi-Fi traffic in the surrounding at different observation points at once. To do so, we leverage an IEEE 802.11 GNU Radio implementation [48] to implement an SDR-based Wi-Fi receiver at SDRs 4A-4B-5A-5B-6A-6B-7A-7B-8A-8B-12-13-14-15-16-20-21-22-23-24 and idle listen on the channel 6 of the Wi-Fi standard. Fig. 19 illustrates the heatmap of the sensed Wi-Fi activity at 20 locations throughout 1-minute long experiments. On average, this experiment reports a higher Wi-Fi activity at certain locations than others, suggesting higher spectrum access opportunities at the latter. This can be due to the proximity to a Wi-Fi access point (AP) or a fairly active Wi-Fi station (STA). Further analysis at higher protocol layers could track these details down and implement intelligent spectrum access techniques.

### 5.5. Ad hoc networks

As mentioned earlier, Arena can implement complex communication schemes such as mesh networking and ad multi-hop communications. We here showcase and implement an instance of WNOS [49] on Arena. WNOS is a wireless network operating system for ad hoc networks, that proposes itself as an intuitive control tool to programmatically dictate the behavior of a network in a simple, centralized way, while this is achieved through distributed control programs running at individual network nodes. WNOS provides automated network control and it interfaces the network controller with a simple control interface. To actuate its distributed control logic, WNOS takes network control problems defined on a centralized abstraction of the network, and automatically generates distributed cross-layer control programs based on distributed optimization theory. These are, then, executed at individual network nodes on an abstract representation of the radio hardware.

To showcase WNOS, we implement a 14-nodes wireless ad-hoc network on Arena, where two source nodes intend to deliver data to two destinations through 12 relay nodes, in a wireless multi-hop fashion. As source, relay, and destination nodes, we use SDRs 1B-3A-3B, SDRs 1A-17-18-2-3-19-20-4A-4B, and SDRs 10–12, respectively. Fig. 20 illustrated the implemented 14-node network deployment on Arena. Then, we employ WNOS to control such a network and seamlessly dictate two different network behaviors, namely *max-rate* and *min-power* (see [49] for details). Finally, we measure the network performance in terms of delivered packets per second at the two destinations. The network performance under the two different control problems is reported in Fig. 20 for a 90 s experiment. The figure illustrates how WNOS is effective in dictating different network behaviors for the same network deployment, enabling automated SDN-like solutions for distributed ad hoc networks. The reader is referred to [49,50] for further details.

## 6. Related work

Over the last decade, open-access experimental SDR-based testbeds have been of paramount importance in testing protocols and technologies both for academic research purposes and for industrial applications. Universities, companies, and consortia have worked on designing a multitude of different experimental testing platforms varying per scale, computational power, deployment environment, and channel characteristics [5,51–63].

Early efforts such as UFMG [58], UFRGS [59], LESC [62], and Iris [61] laid the foundation for larger scale deployments and future design choices. Among others, experimental platforms such as CorteXlab [53] and NITOS Future Internet Facility [54], provided researchers with the tools for indoor Wi-Fi and LTE experimental evaluation, while UNIVBRIS [60] allows outdoor small-city scale testing for the IoT paradigm. Regarding indoor design efforts, we highlight CORNET [64,65], an under-development testbed envisioning 48 nodes deployed across the hallways of a four-story building on campus; and ORBIT [52], a mixed deployment of SDRs and commercial devices ideal for non-line-of-sight experiments and wall penetration testing. Among the other SDR-based testing platforms it is worth mentioning Colosseum [5], and the

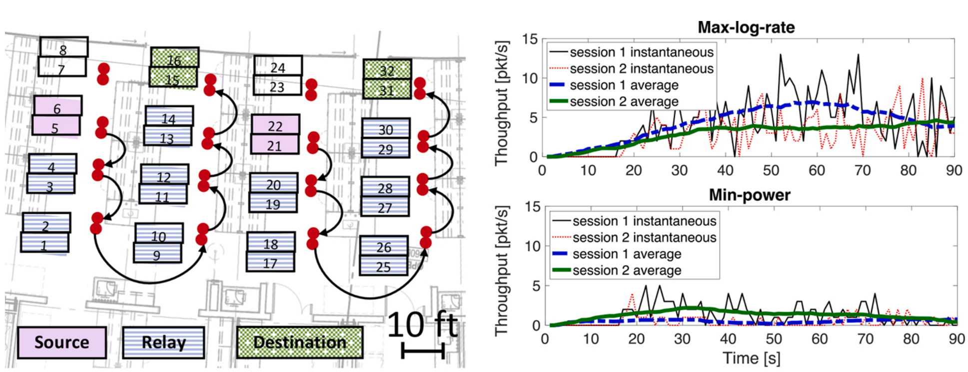


(a) WNOS evaluation topology on Arena. The two sources employ the relays to deliver data to the two destinations in a wireless multi-hop fashion.

(b) WNOS individual and aggregate network throughput performance for two different control problems: max-log-rate (top) and min-power (bottom).

**Fig. 20.** Ad hoc network assessment on Arena.

Drexel Grid [63]. The former is a gargantuan channel emulator tailored to the DARPA Spectrum Collaboration Challenge that allows researchers to test devised approaches on channels emulated through 128 two-channel SDRs. The latter, instead, is a ceiling testbed with some USRPs coexisting with simulated nodes. Both these platforms can use a dataset of recorded real channels characteristics to emulate over-the-air radio communications. Among the large-city-scale outdoor deployments, three leading efforts are the Platform for Open Wireless Data-driven Experimental Research (POWDER) [55], the Cloud Enhanced Open Software Defined Mobile Wireless Testbed for City-Scale Deployment (COSMOS) [56], and the Aerial Experimentation and Research Platform for Advanced Wireelss (AERPAW)[66]. These work-in-progress wireless platforms are based on dozens of nodes supporting different transmission technologies such as cellular networks, Wi-Fi, and drone communications, which allow researchers to evaluate their systems in outdoor city-scale configurations.

Thanks to its three-tier design, with all the SDRs housed in a single rack and the antennas cabled onto the ceiling grid, Arena is, to the best of our knowledge, the only by-design symbol-level synchronized testbed with frequency and phase reference distribution right at the radio-board level. Unlike other indoor testing platforms (ORBIT being the only exception), Arena is based on a grid of antennas mounted on the ceiling of a contiguous indoor environment where all testbed nodes have a full line of sight. This deployment is suitable for numerous line-of-sight applications such as long-distance indoor transmissions and multi-hop and MIMO communications while enabling a large number of possible topology configurations. Different from ORBIT, Arena is deployed in an uncontrolled and diverse office environment, characterized by multiple surrounding objects and materials, rich multipath effect, unpredictable human and object mobility. This aspect makes Arena suitable for performance evaluation of WLAN technologies. To the best of our knowledge, Arena is the world's largest wireless indoor testing platform deployed in an office space. In conclusion, Arena d represents a one-of-a-kind Internet-accessible, open-access research platform with unique full-testbed synchronization, large-scale line-of-sight office deployment, and experimental capabilities over real wireless channels.

## 7. Conclusions

In this article, we presented Arena, an open-access wireless testing platform for sub-6 GHz 5G spectrum research. Arena is a unique indoor testing platform in an office space environment. It is based on 12 computational servers, 24 SDRs, and a total of 64 antennas organized in an $8 \times 8$ grid hanging off the ceiling and covering an overall area of 2240 $ft^2$. We revised its updated layout and architecture design choices, and we presented a detailed description of its hardware and software components. Further, we extensively showcased how Arena can be employed to implement and evaluate complex wireless technologies such as MIMO transmission schemes, 5G cellular networks, Artificial Intelligence and dataset generation, Wi-Fi, Cognitive Radios, and ad hoc networks. We hope that its open-access system, its three-tiered architecture, its full symbol level synchronization, and its unique line of sight indoor ceiling-grid layout ensuring reconfigurable, scalable, and repeatable real-time real wireless channel experiments will foster the development of new 5G-and-beyond technologies, Internet of Things (IoT), Artificial Intelligence powered wireless networks, cognitive radio, dynamic spectrum access, and massive MIMO applications.

## Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was supported in part by the US National Science Foundation under Grant CNS-1618727 and in part by the US Office of Naval Research under Grants N00014-19-1-2409 and N00014-20-1-2132.

## Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.comnet.2020.107436.

## References


[1] P. Cerwall, P. Jonsson, R. Möller, S. Bävertoft, S. Carson, I. Godor, P. Kersch, A. Kälvemark, G. Lemne, P. Lindberg, Ericsson mobility report, On the Pulse of the Networked Society. Hg. v. Ericsson (2015).

[2] D. Evans, The internet of things: how the next evolution of the internet is changing everything, CISCO white paper (2011).

[3] Platforms, for Advanced Wireless Research https://advancedwireless.org.

[4] A. Gosain, Platforms for Advanced Wireless Research: Helping Define a New Edge Computing Paradigm. Proc. of Technologies for the Wireless Edge Workshop, New Delhi, India, 2018, p. 33.

[5] Spectrum, Collaboration Challenge (SC2) https://www.spectrumcollaborationchallenge.com 2020.

[6] Colosseum, at Northeastern University http://www.colosseum.net 2020.

[7] L. Bertizzolo, L. Bonati, E. Demirors, T. Melodia, Arena: A 64-antenna SDR-based Ceiling Grid Testbed for Sub-6 GHz Radio Spectrum Research. Proc. of IEEE Intl. Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization, Los Cabos, Mexico, 2019.

[8] L. Bertizzolo, L. Bonati, E. Demirors, T. Melodia, Demo: Arena: A 64-antenna SDR-based Ceiling Grid Testbed for Sub-6 GHz Radio Spectrum Research. Proc. of IEEE Intl. Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH), Los Cabos, Mexico, 2019.

[9] Ettus, Research https://files.ettus.com/manual/page_calibration.html, 2020.

[10] S. Corum, J.D. Bonior, R.C. Qiu, N. Guo, Z. Hu, Evaluation of phase error in a software-defined radio network testbed. IEEE Southeastcon, Orlando, FL, USA, 2012, pp. 1–4.

[11] Ettus, Research https://kb.ettus.com/Synchronization_and_MIMO_Capability_with_USRP_Devices#Other_Variables_That_Effect_Phase_Alignment, 2020a.

[12] Ettus, Research https://www.ettus.com, 2020b.

[13] AWG, RG8-CMP Datasheet https://rubimages-liberty.netdna-ssl.com/spec/RG8-CMP%20Specification.pdf, 2020.

[14] Wanshih, Electronic Co. https://kb.ettus.com/images/9/9e/ettus_research_vert2450_datasheet.pdf, 2020.

[15] B. Bloessl, M. Segata, C. Sommer, F. Dressler, An IEEE 802.11 a/g/p OFDM Receiver for GNU Radio. Proc. of Workshop on Software Radio Implementation Forum, Hong Kong, China, 2013, pp. 9–16.

[16] I. Gomez-Miguelez, A. Garcia-Saavedra, P. Sutton, P. Serrano, C. Cano, D. Leith, srsLTE: An Open-source Platform for LTE Evolution and Experimentation. Proc. of ACM WiNTECH, New York City, NY, USA, 2016, pp. 25–32.

[17] N. Nikaein, R. Knopp, F. Kaltenberger, L. Gauthier, C. Bonnet, D. Nussbaum, R. Ghaddab, OpenAirInterface: an open LTE network in a PC. Proc. of International Conference on Mobile Computing and Networking, Maui, HI, USA, 2014, pp. 305–308.

[18] F. Kaltenberger, G. de Souza, R. Knopp, H. Wang, The openairinterface 5G New Radio implementation: Current status and roadmap. Proc. of ITG Intl. Workshop on Smart Antennas, Vienna, Austria, 2019, pp. 1–5.

[19] B. Bloessl, C. Leitner, F. Dressler, C. Sommer, A GNU radio-based IEEE 802.15. 4 testbed, GI/ITG KuVS Fachgespräch Drahtlose Sensornetze (FGSN) (2013) 37–40.

[20] M. Knight, B. Seeber, Decoding LoRa: Realizing a modern LPWAN with SDR. Proc. of the GNU Radio Conference 1, Boulder, CO, USA, 2016.

[21] J. Tapparel, O. Afisiadis, P. Mayoraz, A. Balatsoukas-Stimming, A. Burg, An open-Source lora physical layer prototype on GNU radio, arXiv preprint arXiv: 2002.08208 (2020).

[22] B. Bloessl, M. Segata, C. Sommer, F. Dressler, Towards an Open Source IEEE 802.11 p stack: A full SDR-based transceiver in GNU Radio. Proc. of IEEE Vehicular Networking Conference, Boston, MA, USA, 2013, pp. 143–149.IEEE

[23] E.A. Thompson, N. Clem, I. Renninger, T. Loos, Software-defined GPS receiver on USRP-platform, Journal of Network and Computer Applications 35 (4) (2012) 1352–1360.

[24] S. Gunawardena, T. Pany, GNSS SDR Metadata Standard Working Group Report. Proc. of Intl. Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2015), Tampa, FL, USA, 2015, pp. 3218–3221.

[25] G. Girau, A. Tomatis, F. Dovis, P. Mulassano, Efficient software defined radio implementations of GNSS receivers. Proc. of IEEE International Symposium on Circuits and Systems, Baltimore, MD, USA, 2007, pp. 1733–1736.

[26] L. Bertizzolo, E. Demirors, Z. Guan, T. Melodia, CoBeam: Beamforming-based Spectrum Sharing With Zero Cross-Technology Signaling for 5G Wireless Networks. Proc. of IEEE Conference on Computer Communications (INFOCOM), 2020.

[27] The Linux Foundation, Open network automation platform architecture, 2018, (https://tinyurl.com/spyehg9).

[28] O-RAN Alliance, O-RAN: Towards an Open and Smart RAN, 2018, (https://tinyurl.com/t5wlpcz).
[29] S. D'Oro, F. Restuccia, T. Melodia, Toward operator-to-Waveform 5G radio access network slicing, IEEE Commun. Mag. 58 (4) (2020).
[30] S. DOro, F. Restuccia, T. Melodia, S. Palazzo, Low-complexity distributed radio access network slicing: algorithms and experimental results, IEEE/ACM Trans. Networking 26 (6) (2018) 2815–2828.
[31] S. D'Oro, F. Restuccia, A. Talamonti, T. Melodia, The Slice Is Served: Enforcing Radio Access Network Slicing in Virtualized 5G Systems. Proc. of IEEE Conference on Computer Communications (INFOCOM), Paris, France, 2019.
[32] S. D'Oro, L. Bonati, F. Restuccia, M. Polese, M. Zorzi, T. Melodia, Sl-EDGE: Network Slicing at the Edge. Proc. of ACM MobiHoc, 2020.
[33] F. Restuccia, T. Melodia, DeepWiERL: Bringing Deep Reinforcement Learning to the Internet of Self-Adaptive Things. Proc. of IEEE Conference on Computer Communications (INFOCOM), Toronto, ON, Canada, 2020.
[34] F. Restuccia, T. Melodia, Big Data Goes Small: Real-Time Spectrum-Driven Embedded Wireless Networking through Deep Learning in the RF Loop. Proc. of IEEE Conference on Computer Communications (INFOCOM), Paris, France, 2019.
[35] T.J. O'Shea, T. Roy, T.C. Clancy, Over-the-Air deep learning based radio signal classification, IEEE J. Sel. Top. Signal Process. 12 (1) (2018) 168–179.
[36] F. Restuccia, S. D'Oro, A. Al-Shawabka, M. Belgiovine, L. Angioloni, S. Ioannidis, K. Chowdhury, T. Melodia, DeepRadioID: Real-Time Channel-Resilient Optimization of Deep Learning-based Radio Fingerprinting Algorithms. ACM MobiHoc, Catania, Italy, 2019, pp. 51–60.
[37] O. Naparstek, K. Cohen, Deep multi-user reinforcement learning for distributed dynamic spectrum access, IEEE Trans. Wireless Commun. 18 (1) (2019) 310–323.
[38] J. Jagannath, N. Polosky, A. Jagannath, F. Restuccia, T. Melodia, Machine learning for wireless communications in the internet of things: A Comprehensive survey, Ad Hoc Netw. 93 (2019) 101913.
[39] C. Zhang, P. Patras, H. Haddadi, Deep learning in mobile and wireless networking: asurvey, IEEE Communications Surveys & Tutorials 21 (3) (2019) 2224–2287.
[40] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (7553) (2015) 436.
[41] A. Al-Shawabka, F. Restuccia, S. DOro, T. Jian, B.C. Rendon, N. Soltani, J. Dy, S. Ioannidis, K. Chowdhury, T. Melodia, Exposing the Fingerprint: Dissecting the Impact of the Wireless Channel on Radio Fingerprinting. Proc. of IEEE Conference on Computer Communications (INFOCOM), Toronto, ON, Canada, 2020.
[42] R. Francesco, D. Salvatore, M. Tommaso, Securing the internet of things in the age of machine learning and software-Defined networking, IEEE Internet Things J. 5 (6) (2018) 4829–4842.
[43] N.I. Osman, E.B. Abbas, Simulation and modelling of LoRa and sigfox low power wide area network technologies. Proc. of Intl. Conference on Computer, Control, Electrical, and Electronics Engineering (ICCCEEE), Al Gezira, Sudan, 2018, pp. 1–5.IEEE
[44] N. Akhtar, A. Mian, Threat of adversarial attacks on deep learning in computer vision: a survey, IEEE Access 6 (2018) 14410–14430.
[45] R. Zhang, Y. Liang, Exploiting multi-Antennas for opportunistic spectrum sharing in cognitive radio networks, IEEE J. Sel. Top. Signal Process. 2 (1) (2008) 88–102.
[46] I. Akyildiz, W. Lee, M. Vuran, S. Mohanty, A survey on spectrum management in cognitive radio networks, IEEE Commun. Mag. 46 (4) (2008) 40–48.
[47] D. Niyato, E. Hossain, Competitive spectrum sharing in cognitive radio networks: a dynamic game approach, IEEE Trans. on Wireless Communications 7 (7) (2008) 2651–2660.
[48] B. Bloessl, M. Segata, C. Sommer, F. Dressler, Performance assessment of IEEE 802.11p with an open source SDR-based prototype, IEEE Trans. on Mobile Computing 17 (5) (2018) 1162–1175.
[49] Z. Guan, L. Bertizzolo, E. Demirors, T. Melodia, WNOS: An Optimization-based Wireless Network Operating System. Proc. of ACM MobiHoc, Los Angeles, CA, USA, 2018.
[50] Z. Guan, L. Bertizzolo, E. Demirors, T. Melodia, Demo abstract: WNOS: An optimization-based wireless network operating system. Proc. of IEEE Conference on Computer Communications (INFOCOM), Honolulu, HI, USA, 2018.
[51] CORNET, : Cognitive Radio Network Testbed https://cornet.wireless.vt.edu, 2019.
[52] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, M. Singh, Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols. Proc. of IEEE WCNC, New Orleans, LA, USA, 2005.
[53] M. Dardaillon, K. Marquet, T. Risset, A. Scherrer, Software Defined Radio Architecture Survey for Cognitive Testbeds. Proc. of IEEE IWCMC, Limassol, Cyprus, 2012.
[54] K. Pechlivanidou, K. Katsalis, I. Igoumenos, D. Katsaros, T. Korakis, L. Tassiulas, NITOS Testbed: A Cloud based Wireless Experimentation Facility. Proc. of IEEE ITC, Karlskrona, Sweden, 2014.
[55] J. Breen, A. Buffmire, J. Duerig, K. Dutt, E. Eide, M. Hibler, D. Johnson, K. Kumar, L. Sneha, E. Lewis, D. Maas, A. Orange, N. Patwari, D. Reading, R. Ricci, D. Schurig, L.B. Stoller, K. Van der Merwe, K. Webb, G. Wong, POWDER: Platform for Open Wireless Data-driven Experimental Research. Proc. of ACM International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH), 2020.
[56] COSMOS, : Cloud Enhanced Open Software Defined Mobile Wireless Testbed for City-Scale Deployment https://cosmos-lab.org, 2019.
[57] UEFS, Testbed http://futebol.inf.ufes.br, 2019.
[58] UFMG, Testbed http://futebol.dcc.ufmg.br, 2019.
[59] UFRGS, Testbed http://futebol.inf.ufrgs.br, 2019.
[60] UNIVBRIS, Testbed https://www.bristolisopen.com, 2019.
[61] P.D. Sutton, J. Lotze, H. Lahlou, S.A. Fahmy, K.E. Nolan, B. Ozgul, T.W. Rondeau, J. Noguera, L.E. Doyle, Iris: an architecture for cognitive radio networking testbeds, IEEE Commun. Mag. 48 (9) (2010) 114–122.
[62] LESC, CR/SDR Testbeds http://lesc.det.unifi.it/en/node/194, 2019.
[63] K.R. Dandekar, S. Begashaw, M. Jacovic, A. Lackpour, I. Rasheed, X. Rivas Rey, C. Sahin, S. Shaher, G. Mainland, Grid software defined radio network testbed for hybrid measurement and emulation. Proc. of IEEE SECON, Boston, MA, USA, 2019.
[64] N. Sharakhov, V. Marojevic, F. Romano, N. Polys, C. Dietrich, Visualizing real-time radio spectrum access with CORNET3D. Proc. of ACM Web3D, Vancouver, BC, Canada, 2014.
[65] S. Kikamaze, V. Marojevic, C. Dietrich, Demo: Spectrum access system on cognitive radio network testbed. Proc. of ACM WiNTECH, Snowbird, UT, USA, 2017.
[66] AERPAW, : Aerial Experimentation and Research Platform for Advanced Wirelelss https://aerpaw.org/, 2020.


**Lorenzo Bertizzolo**is a candidate for Ph.D. in Computer Engineering and research assistant at the Institute for the Wireless Internet of Things at Northeastern University. He earned his B. S. and his M.S. with honors in Computer and Communication Networks Engineering from Politecnico di Torino, Italy in 2014 and 2015, respectively. His research focus is software-defined control for wireless networks, MIMO systems, and Wireless Networks of Drones. He is also a collaborator of AT&T Labs Research, working on the integration of Unmanned Aerial System into the next generations' cellular networks.

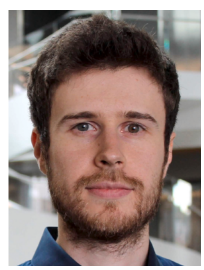

**Leonardo Bonati**received his B.S. in Information Engineering and his M.S. in Telecommunication Engineering from University of Padova, Italy in 2014 and 2016, respectively. He is currently pursuing a Ph.D. degree in Computer Engineering at Northeastern University, MA, USA. His research interests focus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

**Emrecan Demirors** is a Research Assistant Professor with the Department of Electrical and Computer Engineering at Northeastern University. He is conducting research at the Wireless Networks and Embedded Systems Laboratory. Previously, he was an Associate Research Scientist with the Department of Electrical and Computer Engineering at Northeastern University, from 2017 to 2019. He received my Ph.D. degree in Electrical and Computer Engineering from Northeastern University in 2017, under the supervision of Professor Tommaso Melodia. He had previously received my B.S. and M.S degrees in Electrical and Electronics Engineering from Bilkent University, Ankara, Turkey in 2009 and 2011, respectively, under the supervision of Professor Hayrettin Koymen. From 2010 to 2011, he was a Systems Engineer at Meteksan Defence Industry Inc., Ankara, Turkey.

**Amani Al-shawabka** is currently a Ph.D. candidate at Northeastern University in the Electrical and Computer Engineering department. She received her Master's degree in Computer Engineering from Northeastern University in 2019. She started her professional career in mobile network companies (Orange and Batelco groups) after receiving her B.S. degree in Communication Engineering from Yarmouk University, Jordan. Her research interests are wireless networks, wireless fingerprinting and embedded systems and network security.

**Salvatore D'Oro** received received his Ph.D. degree from the University of Catania in 2015. He is currently an Associate Research Scientist at Northeastern University. In 2015, 2016 and 2017 he organized the 1st, 2nd and 3rd Workshops on COmpetitive and COoperative Approaches for 5G networks (COCOA). He also served on the Technical Program Committee (TPC) of the IEEE Conference on Standards for Communications and Networking (CSCN'18), Med-Hoc-Net 2018 and the CoCoNet8 workshop at IEEE ICC 2016. He serves on the TPC of Elsevier Computer Communications journal. Dr. D'Oro is also a reviewer for major IEEE and ACM journals and conferences. Dr. D'Oro's research interests include game-theory, optimization, learning and their applications to telecommunication networks. He is a Member of the IEEE.

**Francesco Restuccia** received his Ph.D. in Computer Science from Missouri University of Science and Technology, Rolla, MO, USA, in 2016, and his M.S. and B.S. in Computer Science and Engineering with highest honors from the University of Pisa, Italy in 2011 and 2009, respectively. Currently, he is an Associate Research Scientist with the Institute for the Wireless Internet of Things, Department of Electrical and Computer Engineering, Northeastern University, Boston, MA, USA. His research interests lie in the modeling, analysis, and experimental evaluation of wireless networked systems, with applications to pervasive computing and the Internet of Things. Dr. Restuccia has published over 25 papers in top venues, as well as coauthoring 8 pending US patents and 2 book chapters. He regularly serves as a TPC Member for conferences such as IEEE INFOCOM and ACM MobiHoc, and is a reviewer for several ACM and IEEE conferences and journals. Dr. Restuccia is the recipient of the 2019 ISSNAF Mario Gerla Award for Young Investigators in Computer Science. He is a Member of the IEEE and the ACM

**Tommaso Melodia**is the William Lincoln Smith 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 is the Editor in Chief of Computer Networks. 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.