Challenge: COSMOS: A City-Scale Programmable Testbed for Experimentation with Advanced Wireless

Dipankar Raychaudhuri[1], Ivan Seskar[1], Gil Zussman[2], Thanasis Korakis[3], Dan Kilper[4],
Tingjun Chen[2], Jakub Kolodziejski[1], Michael Sherman[1], Zoran Kostic[2], Xiaoxiong Gu[5],
Harish Krishnaswamy[2], Sumit Maheshwari[1], Panagiotis Skrimponis[3], Craig Gutterman[2]
[1]WINLAB, Rutgers University, [2]Electrical Engineering, Columbia University,
[3]Electrical and Computer Engineering, NYU, [4]College of Optical Sciences, University of Arizona, [5]IBM Research
{ray, seskar}@winlab.rutgers.edu, {gil, tingjun}@ee.columbia.edu,
korakis@nyu.edu, dkilper@optics.arizona.edu, xgu@us.ibm.com
`https://cosmos-lab.org/`

ABSTRACT

This paper focuses on COSMOS – **C**loud enhanced **O**pen **S**oftware defined **MO**bile wireless *testbed* for city-**S**cale deployment. The COSMOS testbed is being deployed in West Harlem (New York City) as part of the NSF Platforms for Advanced Wireless Research (PAWR) program. It will enable researchers to explore the technology "sweet spot" of ultra-high bandwidth and ultra-low latency in the most demanding real-world environment. We describe the testbed's architecture, the design and deployment challenges, and the experience gained during the design and pilot deployment. Specifically, we describe COSMOS' computing and network architectures, the critical building blocks, and its programmability at different layers. The building blocks include software-defined radios, 28 GHz millimeter-wave phased array modules, optical transport network, core and edge cloud, and control and management software. We describe COSMOS' deployment phases in a dense urban environment, the research areas that could be studied in the testbed, and specific example experiments. Finally, we discuss our experience with using COSMOS as an educational tool.

CCS CONCEPTS

• **Networks** → **Network architectures**; *Wireless access networks*; *Wireless access points, base stations and infrastructure*; *Network experimentation*; • **Hardware** → *Wireless devices*;

KEYWORDS

Testbed, experiments, wireless networks and communications, software-defined radios, millimeter-wave, 5G, edge cloud, optical x-haul, programmability, STEM education

ACM Reference Format:
Dipankar Raychaudhuri, Ivan Seskar, Gil Zussman, Thanasis Korakis, Dan Kilper, Tingjun Chen, Jakub Kolodziejski, Michael Sherman, Zoran Kostic, Xiaoxiong Gu, Harish Krishnaswamy, Sumit Maheshwari, Panagiotis Skrimponis, Craig Gutterman. 2020. Challenge: COSMOS: A City-Scale Programmable Testbed for Experimentation with Advanced Wireless . In *The 26th Annual International Conference on Mobile Computing and Networking (MobiCom '20), September 21–25, 2020, London, United Kingdom.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3372224.3380891

1 INTRODUCTION

Wireless technology has evolved at a remarkable rate, with data rates increasing by four orders of magnitude over the past twenty years [1–4]. The push towards 5G will continue this trend, with radio access links soon operating at 1 Gbps or higher [5–7]. Recent efforts focus on reducing access network latency from 10s of milliseconds to the order of 1 ms [8, 9].

New wireless communication technologies such as millimeter-wave (mmWave), large/distributed MIMO, and dynamic spectrum access will enable these order-of-magnitude gains and usher in a new era of applications. However, *a major challenge is the lack of a programmable testbed for allowing the academic research community as well as the wireless industry to explore in the most demanding real-world environments the technology "sweet spot" of ultra-high bandwidth and ultra-low latency.* Such a capability is required to enable a broad new class of real-time applications including augmented/virtual reality (AR/VR) [10–15] and cloud-based autonomous vehicles [16, 17]. It has been recently realized that enabling the transition from academic research to the wireless industry requires realistic city-scale testbeds [18, 19].

Hence, this paper focuses on the COSMOS testbed [20] which is designed and being deployed in order to address this challenge. The testbed is being deployed in West Harlem (New York City) as part of the NSF Platforms for Advanced Wireless Research (PAWR) program by Rutgers University, Columbia University, and NYU in partnership with New York City (NYC), City College of New York (CCNY), IBM, University of Arizona, and Silicon Harlem. This paper outlines the challenges associated with the design, development, and deployment of the testbed as well as the experience gained during these phases. Specifically, the realization of ultra-high bandwidth and ultra-low latency wireless applications involves research

---

MobiCom '20, September 21–25, 2020, London, United Kingdom

ACM ISBN 978-1-4503-7085-1/20/09…$15.00
https://doi.org/10.1145/3372224.3380891

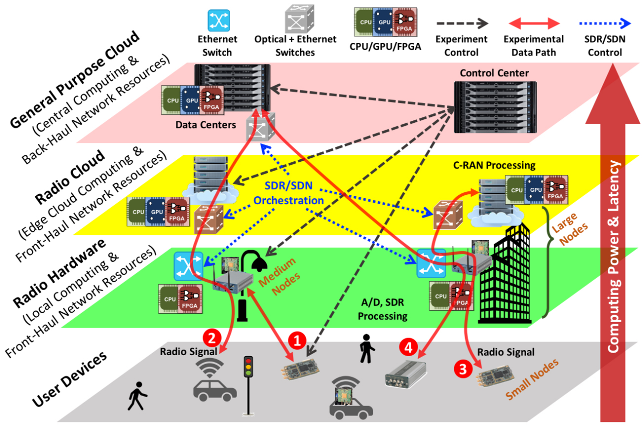

Figure 1: **COSMOS' multi-layered computing architecture, where paths 1–4 demonstrate the data paths of example experiments with local/remote computing.**

not only on faster radio links, but also on the system as a whole, including aspects such as spectrum use, networking, and edge cloud and computing.

Therefore, there is a need to design a testbed that will enable researchers to conduct accurate experiments over a broad range of new system designs. It should incorporate emerging spectrum techniques (e.g., dynamic spectrum sharing [21–23], mmWave [24–32], and full-duplex [33–38]), techniques for Gbps+ radio access such as large/distributed MIMO [39–49], low latency mobile core network protocols [50, 51], and edge cloud services [52–59]. A testbed focusing on high bandwidth and low latency will also provide the capabilities needed for Internet-of-Things (IoT) [60], e-Health [61], and smart communities [62] applications.

Achieving this multi-dimensional goal requires careful consideration of the testbed's design and deployment. Therefore, we first outline COSMOS' multi-layered computing and network architecture, where the testbed relies on *open-source software* and *allows for different levels of programmability and control* from the user devices to the cloud (see Fig. 1). The testbed's 3-phase envisioned deployment covers a 1 sq. mile dense urban area and includes 9 rooftop-installed *large* nodes, ~40 building-side or lightpole-mounted *medium* nodes, and ~200 *small* fixed or mobile nodes, which are the testbed "users". It also includes various computing resources, optical equipment, and an underlying fiber network.

We describe the design challenges and programmability aspects related to the following building blocks: (i) software-defined radios (SDRs) and the IBM 28 GHz programmable phased array antenna modules (at the nodes), (ii) optical transport network, and (iii) core and edge cloud. We also discuss the software-defined networking (SDN) framework and the control and management software that enable experimentation with different radio nodes and heterogeneous cloud computing resources (i.e., CPUs, GPUs, and FPGAs) in the COSMOS testbed.

To provide experimenters with sufficient scale and realism, *COSMOS is being deployed in West Harlem, a representative dense urban environment, with diverse radio propagation and pedestrian/vehicle/traffic characteristics*. Such deployment requires the support of NYC and the local community and poses unique challenges not commonly addressed by the research community. Specifically, the deployment of COSMOS consists of 3 phases and involves the facilities of 4 universities (Rutgers, Columbia, NYU, and CCNY) and is based on extensive support from several city agencies and the local community. We describe the experience gained in the Pilot Phase deployment (completed in May 2019). We outline the planned deployment in Phase 1 (expected by the end of 2020) and the vision and plan for Phase 2. We also briefly discuss the designation of the testbed area as an FCC innovation zone [63].

We then outline potential research areas that could be studied in COSMOS at different layers of the protocol stack (from the PHY layer to the application layer) and in different network domains (e.g., wireless, wired, and optical). We briefly describe 4 example experiments conducted by the COSMOS team to drive and validate various testbed capabilities. These examples can help researchers envision and plan their own experiments.

Finally, we discuss activities focusing on transforming COSMOS into an innovative learning platform for middle and high students that would help bridge the digital divide and provide significant educational benefits for the local community. Specifically, we discuss the development of the COSMOS Education Toolkit [64] that is already used in middle and high schools.

Throughout the paper, we outline the significant challenges related to designing and deploying the COSMOS testbed in a real urban setting. These include the following:

1. **Deployment Challenges**: The envisioned deployment of the testbed will cover an area of about 1 sq. mile. The deployment will rely on collaboration with several NYC agencies, various local community stakeholders, the infrastructure of 4 universities, and contributions of 30 PAWR industry consortium members. Managing these collaborations and leveraging these contributions while achieving the scientific and research goals of the research community is challenging.
2. **Management Challenges**: The control and management of such a large scale testbed in a dense urban environment with 24/7 operation, without causing interference to other systems, and with timely technical support is extremely challenging.
3. **Technology Development and Experimentation Challenges**: COSMOS' key technological building blocks include multiple emerging technologies whose development and integration as well as support for experimenter-friendly operation present challenges.

Accordingly, the rest of the paper is organized as follows. We review related work in Section 2 and discuss the testbed design and architecture in Section 3. We present the technology building blocks in Section 4 and discuss the envisioned deployment plan in Section 5. Section 6 describes the research areas that could be studied in the testbed, and reviews specific example experiments. We discuss the outreach and education plan in Section 7, and conclude in Section 8.

2 RELATED WORK

Several testbeds with multi-technology capabilities aimed to address the new challenges posed by 5G and future wireless networks have been recently developed and deployed. For example, Bristol Is Open [65] is a programmable testbed deployed in Bristol, UK, that provides a smart city development and research platform for wireless network connectivity, IoT hosting, and cloud computing. The ADRENALINE [66] testbed deployed in the Catalania region, Spain, is a circuit-switched optical testbed designed for experimental research on large-scale optical transport networks and distributed edge computing. POWDER-RENEW [67] is another advanced wireless testbed that shares some design/operation considerations with COSMOS, and is currently under development. It focuses on massive MIMO wireless and will support wireless experiments in Salt Lake City by using the cloud computing resources provided by CloudLab [68].

In comparison, COSMOS is deployed in a densely populated urban environment in West Harlem (NYC), with unique technological emphasis on mmWave wireless, converged optical-wireless networking, edge cloud, and full programmability. The architecture of COSMOS is informed by networking and wireless testbeds such as GENI, Emulab, PlanetLab, OneLab, CloudLab, CIAN TOAN [68–73], and most notably ORBIT [74], which is a large-scale indoor emulation testbed with similar design requirements and scale. However, developing and deploying a next-generation advanced wireless testbed at city-scale in a real dense urban setting poses unique challenges that will be reviewed in this paper.

3 TESTBED DESIGN & ARCHITECTURE

3.1 Multi-Layered Computing Architecture

To provide the required capabilities to experimenters, COSMOS must address new requirements including support for mmWave communications, wideband radio signal processing (500 MHz or more), effective virtualization of radio resources, low latency front- and back-haul, tightly coupled edge cloud, and real-world city-scale deployment. These requirements pose significant challenges.

To address these challenges, COSMOS' architecture is based on a *multi-layered computing system* (see Fig. 1). It is built in a bottom-up manner with commodity components, programmable hardware, and open-source software. Three types of hardware components are included: (i) SDR nodes (*user devices* and *radio hardware*), (ii) edge cloud servers (*radio cloud*), and (iii) *general purpose cloud*. In particular, COSMOS includes three types of radio nodes (see Section 4): *large* (rooftop installed wide-area base stations [BSs]), *medium* (building side or lightpole-mounted micro-cellular BSs), and *small* (fixed or mobile).

COSMOS also integrates edge cloud technology, including commodity CPUs/GPUs/FPGAs, for achieving computing speeds needed to support cloud radio access networks (C-RANs), network function virtualization (NFV), and low-latency cloud applications. Signal processing and NFV can be flexibly partitioned between a local SDR with CPU/FPGA assists (i.e., radio hardware), and a remote C-RAN with CPU/GPU/FPGA assists (i.e., radio cloud). These two computing layers are backed up by the general purpose cloud layer that serves network- and application-layer functions associated with an experiment.

The remote accessibility and open programmability of COSMOS (see Section 4) allows users to orchestrate their experiments involving various wired/wireless resources by flexibly configuring the network topologies and computing chains. For instance, Fig. 1 shows the data paths used in example experiments supported by COSMOS' multi-layered computing architecture (see Section 6 for the detailed experiments): (1) local computing by a medium node, (2) remote computing for a medium node at a nearby large node's servers or at the central computing servers, (3) local computing by large node's computing servers, and (4) remote computing for a large node at the central computing servers. Note that COSMOS can also support a broader set of experiments not limited to the four paths mentioned above (e.g., local computing by a large node using its CPU/FPGA).

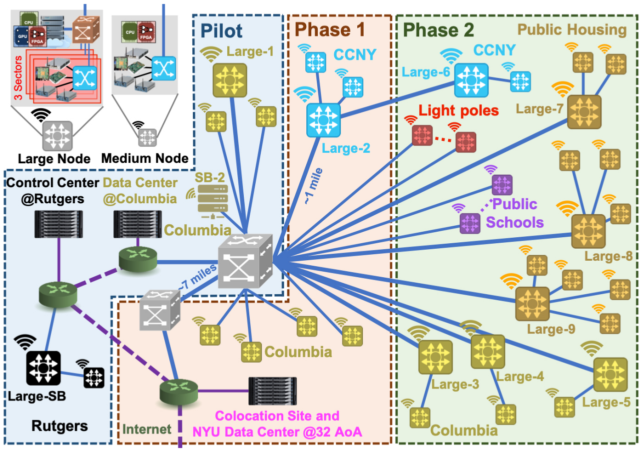

Figure 2: **COSMOS' network infrastructure, connectivity, and envisioned deployment phases. Components include: large and medium nodes whose colors correspond to the deployment sites, sandboxes (SBs) at Rutgers and Columbia, core optical switching at Columbia (large gray cube), optical switching at the remote data center of NYU in the 32 AoA colocation site (small gray cube), compute at Columbia, NYU data center and Rutgers, optical switching within a large node (brown cube), and the underlying optical transport network (see Section 4.3 and Fig. 7 for details).**

3.2 Network Infrastructure & Connectivity

To keep pace with the significantly increased wireless link bandwidth and to effectively integrate the emerging C-RANs, COSMOS is designed to incorporate a fast programmable core network for providing connections across different computing layers. This core network consists of mostly 100 Gbps+ fiber, as well as free-space optical (FSO) and microwave backhaul technologies interconnected with an SDN switching fabric for both minimum latency and flexibility in setting up experimental network topologies. COSMOS integrates advanced optical switching technology based on wavelength-division multiplexing (WDM) switch fabrics and radio over fiber (RoF) interfaces to achieve ultra-low latency connections to edge and central clouds.

As illustrated in Fig. 2, the COSMOS testbed includes sites at 4 universities: Rutgers, Columbia, NYU, and CCNY. In particular, Columbia (both the Morningside and Manhattanville campuses) and CCNY are located at the boundary of the targeted West Harlem deployment area of ~1 sq. mile (see also Fig. 10). NYU is located in Manhattan and Brooklyn with its data center housed at a colocation site at 32 Avenue of the Americas (32 AoA), which is in Manhattan and is ~7 miles away from the Columbia Morningside campus.

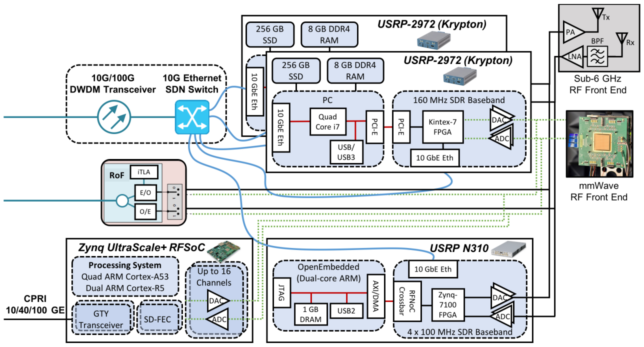

Figure 3: **Example block diagram of a medium node or a sector of a large node, where different node configurations include a *subset* of the major components (e.g., sub-6/28 GHz RF front ends, FPGA-assisted SDRs, and radio over fiber [RoF] interface).**

COSMOS is designed to provide an ultra-low latency interconnect between the large/medium nodes and the edge/central cloud resources. Fig. 2 shows COSMOS' customized optical front-haul network, where large and medium nodes (to be deployed at Columbia, CCNY, and NYC assets in West Harlem) will be connected to the core optical switching in the data center at Columbia. Although most connections are fiber-based in a tree-topology, FSO and microwave backhaul can be opportunistically used between a subset of large and medium nodes. As we will describe in Section 4.3, based on different experimental requirements, a variety of x-haul optical and networking topologies can be overlayed on the fiber plant using SDN.

The optical network connects to two computing racks located at Columbia and the NYU data center (at the 32 AoA collocation site). The control center (located at Rutgers) is connected to the testbed over the Internet, and includes servers for hosting the user portal, scheduling/virtualization services, experiment management services, data repository, etc. Fig. 2 also outlines the phases of the envisioned COSMOS' deployment plan, whose associated challenges (e.g., nodes will be densely deployed in Columbia residential buildings, NYC public housing/schools, and other NYC assets, covering an area of ~1 sq. mile) are described in Section 5.

4 TECHNOLOGY BUILDING BLOCKS

We now describe COSMOS' key building blocks that have been designed to cope with the challenges outlined above.

4.1 Software-Defined Radio (SDR) Nodes

The COSMOS SDR nodes in small/medium/large form factors include the same general components (e.g., CPU, FPGA-assisted SDRs, and antennas). They differ by the number of antennas supported, RF bands/bandwidth covered, physical size, and the level of SDR processing provided. Fig. 3 shows an example block diagram of an SDR node, where different node configurations include a *subset* of the components. To support various experimental capabilities required by the research community, all the radio nodes offer full flexibility of spectrum use in 400 MHz–6 GHz with a subset of the nodes equipped with 28 GHz mmWave capability (see Section 4.2). The nodes also incorporate RF sensing and measurement capabilities to support spectrum use and propagation studies, and operate in conjunction with the edge/central cloud to support different levels of remote processing.

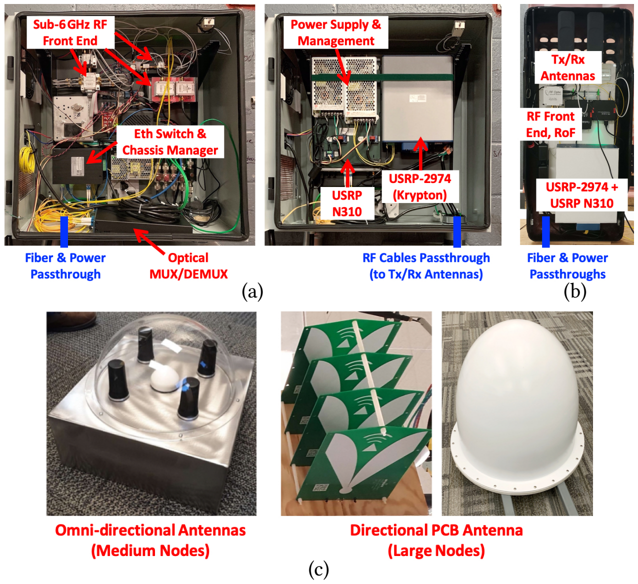

Figure 4: **(a) Developed and deployed *medium* node (or a sector of a *large* node) with the components mounted on the top shelf (left) and bottom (right) shelf, (b) prototype of the light version of the *medium* node for lightpole mounting (*medium-light*) with the components layout, and (c) sub-6 GHz wideband antennas used for the medium (left) and large (right) node.**

Below, we present the COSMOS radio nodes in three form factors, and the core components of sub-6 GHz transceivers.

Small Node. The *small* nodes are portable and tailored for installation in a vehicle or carried in a backpack or a cart, and can be used as fixed or mobile clients and for monitoring functions such as spectrum measurement. The hand-held USRP E312 is the most compact. Alternatively, a near-portable platform based on the Intel NUC and USRP B210/B205mini-i or a larger variant based on a standard ITX motherboard can provide improved performance.

Medium Node. The *medium* nodes are suitable for street-level building-side and lightpole mounting. They can serve as infrastructure micro-cellular BSs or access points (APs). Fig. 4(a) shows a deployed medium node with major modules including USRP-2974/N310, RF front end, Ethernet switches, and fiber connectors (see also Fig. 3). We adopt three configurations based on the support for the 28 GHz phased arrays (see Section 4.2) with different bandwidth requirements:

1. Without mmWave capability: one USRP-2974;
2. With mmWave capability (up to 160 MHz bandwidth): two USRP-2974s with one dedicated to mmWave;
3. With mmWave capability (500 MHz or higher bandwidth): one USRP-2974 and one Zynq UltraScale+ RFSoC platform, which will be dedicated to mmWave.

Each medium node has a separate antenna package with 4 wideband omni-directional antennas, as shown in Fig. 4(c). A light version of the medium node (*medium-light*) with RoF interface (see Fig. 4(b)) will be used for lightpole mounting.

Large Node. The *large* nodes are designed to realize higher powered BSs for wide-area coverage, and are typically deployed on

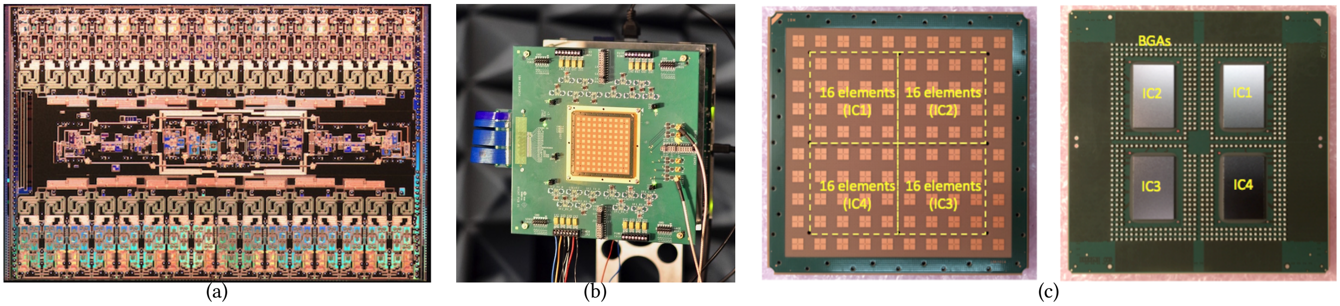

(a) (b) (c)

Figure 5: **(a) The die photograph of the IBM 28 GHz phased array IC implemented in 130 nm SiGe BiCMOS technology, (b) the silicon-based 28 GHz phased array antenna module (PAAM) mounted on a test board, and (c) top and bottom views of a fully assembled 28 GHz PAAM [75].**

rooftops at 100 feet or higher. A large node consists of *three sectors*, each including identical components as a medium node (with the same three design configurations), and 4 wideband directional antennas, as shown in Fig. 4(c). The three sectors are connected to the large node's servers through optical and Ethernet switches.

Sub-6 GHz Transceiver. The core of sub-6 GHz transceiver is based on two types of COTS SDRs. The smaller units (e.g., USRP E312) are associated with reduced baseband bandwidth of up to 50 MHz using USB3 interfaces. The more powerful units, e.g., USRP-2974 as shown in Fig. 3, have dual RF daughterboards with 160 MHz baseband and large Xilinx FPGAs.

Since these SDRs span the entire sub-6 GHz band with RF front end complex consisting of power amplifiers (PAs), low-noise amplifiers (LNAs), optional duplexers, etc., the antennas have to be modular to provide efficient coverage. The large/medium nodes have a number of these front ends to support experimental requirements (e.g., full-duplex wireless [36, 38], 900 MHz and 2.4–2.5/2.5–2.7/3.7–4.2/5.1–5.9 GHz). Antennas are selected using an RF switching complex, supporting up to 8 RF front end modules per USRP daughterboard. The RF switching and front ends are controlled by the SDR through the GPIO interface or through an USB/Ethernet-based remote controller.

Programmability and Control. COSMOS' architecture includes CPUs/GPUs/FPGAs that are distributed throughout the testbed (see Fig. 1 and Section 4.4). Depending on experimental requirements, running experiments involves setting up the SDRs and the corresponding computing chains through the control and management software (see Section 4.5). The USRP-2974/N310 SDRs can be controlled from the embedded PC or from a compute node/server, both of which can be loaded with different prepared/customized images. All SDRs can also be programmed in the FPGA scope (e.g., through the open-source RFNoC tool). For example, users can deploy SDR-based open-source OpenAirInterface (OAI) LTE experiments orchestrated by Open Source Mano (OSM) using COSMOS' distributed computing resources.

The frequency bands and maximum transmit power levels of the sub-6 GHz transceivers are automatically configured by COSMOS' control and management software upon testbed reservation, during which the experimental requests submitted by users are transformed into configuration commands to be sent to the nodes. Moreover, the control and management software ensures no collision in the reserved frequency bands when the testbed is simultaneously shared between multiple users (see Section 4.5). Coarse calibration of the SDRs will be performed by the COSMOS team on a regular basis, and finer-grained calibration can be performed by the users.

4.2 28 GHz Phased Array Antenna Module

Another key building block of COSMOS is the IBM 28 GHz phased array antenna module (PAAM) [75], which can enable unique mmWave experimentation at the PHY, link, and network layers with multi-beam support and agile beam steering capability, which is currently not supported by other testbeds. The PAAM employs phased array ICs (see Fig. 5(a)) fabricated in 130 nm SiGe BiCMOS technology. The ICs are assembled with an antenna-in-package array to form the PAAM (see Figs. 5(b) and 5(c)), which was extensively tested and characterized in lab environments.

The PAAMs include complete radio front end functionality (e.g., PAs, LNAs, mixers, and orthogonal phase and amplitude control per element). Each IC includes 32 TRx phase shifting front end elements and features concurrent independent beams in two polarizations in either Tx or Rx operation. Each PAAM includes 4 such ICs and 64 dual-polarized antennas that provide eight 16-element or two 64-element concurrent beams with complex modulation formats (e.g., 256QAM). We adopt two design configurations for integration with a subset of COSMOS' radio nodes:

1. *Two to eight 3 GHz IF interfaces*: This design would enable direct coupling to USRP-2974 which can perform IQ down conversion of the 3 GHz IF;
2. *Two baseband interfaces (one for each polarization)*: This design would include 3 GHz IQ upconverters and downconverters on the board (e.g., from Analog Devices) to enable coupling to FPGA-based PHY layer implementations where the signal bandwidth would not be limited by the USRPs. In particular, the high-performance Zynq UltraScale+ RFSoC platform (see Fig. 3) with 16 channels of RF ADC and RF DAC will be used to implement a complete SDR.

Moreover, the baseband will be developed in phases with the initial phase targeting a minimal configuration for 28 GHz bands, which are likely to be allocated in 100 MHz component carriers with a maximum allocation of 4 carriers. Once the PAAMs are integrated in COSMOS, their multi-beam support and agile beam steering capability can enable experimentation with unique mmWave wireless links and networks based on hybrid beamforming and MIMO. We have also been conducting extensive 28 GHz channel measurements in the COSMOS testbed deployment area in order to characterize the

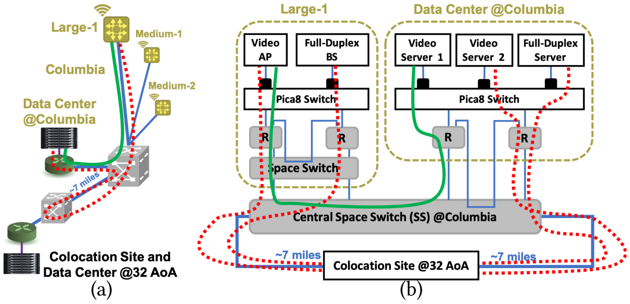

Figure 6: **(a) Example of optical x-haul network topologies that can be created by COSMOS' optical transport network with SDN control, and (b) the optical paths in an example optical-wireless x-haul experiment described in Section 6.**

coverage performance and achievable data rates with preliminary results summarized in [76].

Programmability and Control. The PAAM is controlled by a Zynq-7000 SoC device using a customized API via the USB interface. The following settings can be programmed by the experimenters: (i) number of ICs and number of elements in each IC to be activated, (ii) beamforming mode (e.g., Tx/Rx beamforming in H/V polarization), and (iii) beam steering directions and beamforming weights (amplitude and phase).

4.3 Optical Transport Network

A critical building block of COSMOS is the optical transport network and its unique ability to emulate optical front/mid/backhaul (x-haul) networks, which is currently not supported by other testbeds. It is used to enable high bandwidth and low latency networking and to support experimentation with emerging and future wireless technologies and applications. Using WDM with wavelength switching in colorless, reconfigurable optical add-drop multiplexers (ROADMs)[1], COSMOS enables a wide range of x-haul networks with different topologies in a real metropolitan setting.

Fig. 6(a) demonstrates the motivation for the optical architecture by showing a case in which COSMOS' core switching architecture can create customized networking topologies to emulate a variety of distances for C-RAN remote computing. For example, a large node (Large–1 at Columbia) can perform computing in the data center at Columbia through a short optical route (green, with ~0.3 mile distance), or through a long optical route using dark fiber provided by ZenFi and NYC to the colocation site (red, with ~14 miles round-trip distance). An example experiment demonstrating these scenarios is presented in Section 6. Similarly, it is possible to create various mesh topologies overlayed on the actual fiber plant.

We now provide a high-level description of COSMOS' optical transport network (a technical description appears in [77]). The optical core network and connection to other testbed components are shown in Fig. 7 (the gray cube corresponds to the large gray cube in Fig. 2). In particular, a central Calient S320 320×320 space switch (SS)[2] in the Data Center at Columbia provides a remotely configurable fiber plant network core. Each fiber supports 96 fully

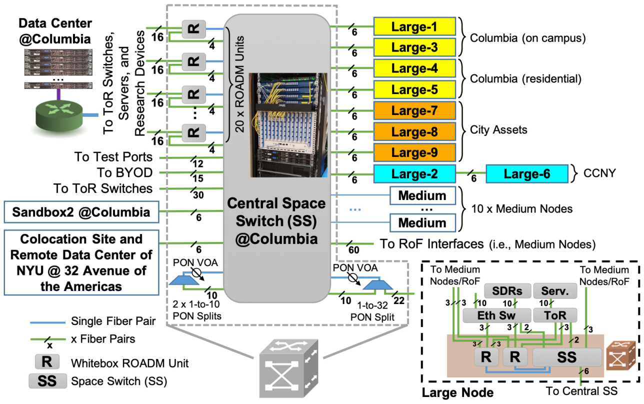

Figure 7: **COSMOS' core optical switching architecture and the switching architecture of a large node.**

transparent channels, which are flexgrid configurable to provision 10/100 Gbps wavelengths. 20 ROADM units are connected to the central SS via 20×4 add/drop fiber pairs. In addition, 20×16 add/drop pairs connect the ROADM units to top of rack (ToR) switches and servers in the data center at Columbia. The 20 fiber pairs per ROADM can be reconfigured to support various requirements from connected servers through top-layer user applications orchestrated by an SDN controller. The central SS is directly connected via dark fiber (provided by ZenFi and NYC) to the colocation site at 32 AoA that also houses the NYU data center.

As indicated by the switching architecture of a large node in Fig. 7 (see also Fig. 2), each large node is equipped with a 16×16 SS and 2 ROADM units (brown cube), Ethernet switches, and ToR switches. 6 fiber pairs connect the central SS and the SS in each large node allowing for a combination of point to point, passive optical network (PON), and ROADM/WDM networks. A subset of the large/medium nodes will also be equipped with RoF capability (see Fig. 3).

Programmability and Control. The optical equipment can be controlled using NETCONF tools. However, due to the sensitivity and safety concerns of the optical equipment (i.e., the SSs, ROADM units, and optical transceivers), a validated API is provided by the COSMOS team for configuring the optical transport network. In particular, the API is implemented using an open-source NETCONF-compatible SDN control plane with customized command line scripts and a Ryu OpenFlow controller (for more details see [78]). It also includes a sufficient number of optical network topologies involving different large/medium nodes and computing resources that can be selected by the users based on experimental requirements. Special full access to the optical equipment may be given to experienced users.

4.4 Core and Edge Cloud

Fig. 8 shows COSMOS' cloud architecture including both the core and edge cloud. The edge cloud computing sites contain three types of computing resources: CPUs, GPUs, and FPGAs in order to provide flexible and powerful signal processing capabilities, while also supporting general purpose computing tasks required for different applications. Software tools are provided to support parallel execution on combinations of CPU/GPU/FPGA. The software also

---

[1] A ROADM enables add, block, pass, or redirect of wavelengths to remotely provision and reconfigure traffic for a WDM network.

[2] A space switch (SS) is used to enable an interconnection between multiple optical input and output ports.

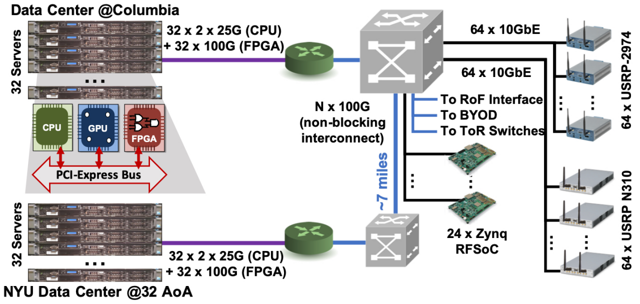

Figure 8: **Logical topology of COSMOS' cloud architecture overlayed on its optical topology, with aggregated link capacities indicated.**

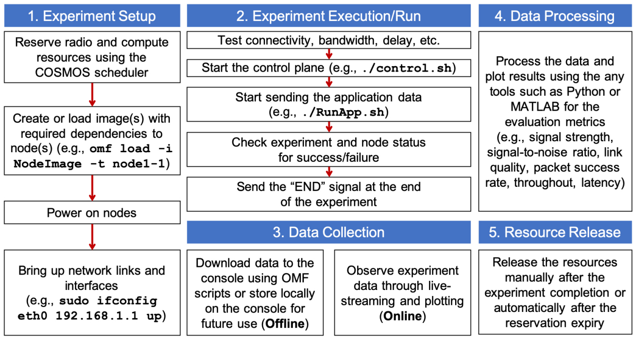

Figure 9: **The procedure for executing an experiment in COSMOS.**

allows convenient software/services development and deployment on a virtualized platform, and access to the "bare-metal" for developing applications/services that need to take full advantage of the ultra-high bandwidth and ultra-low latency of the testbed.

COSMOS' SDN framework integrates native and agent-based control of both wired and wireless resources. This allows SDN experiments to implement application-driven control of optical and data networking functionalities, and radio resources. Further, the SDN framework supports virtualization and allows for logical separation of the same radio or network resource into multiple distinct networks with their own topology and routing protocol [79].

Programmability and Control. The SDN framework supports programmability at L1/L2/L3 (optical/Ethernet/IP), and users have full access to ToR (Ethernet) switches. All the heterogeneous cloud computing resources (CPUs/GPUs/FPGAs, as well as a mix of these) can be accessed and programmed by the users for various wired and wireless experiments. The open-source open network operating system (ONOS) platform and Ryu OpenFlow controllers will be used as standard platforms for SDN and NFV experimentation. Recent multi-tenant controllers will also be used to allow different experiments to run unique SDN controllers simultaneously [80].

4.5 Control and Management Software

Due to the deployment in a dense urban environment (see Section 5) and various supported technologies and experiments (see Section 6), it is more challenging to manage COSMOS than legacy testbeds. The COSMOS control and management software, whose architecture is based on a central controller (similar in concept to SDN), is one of the critical technologies for successful operation of the testbed. In particular, it interfaces with experimenters, reserves/schedules and sets up experimental resources, and manages experiment execution. COSMOS leverages and enhances the OMF testbed control framework [81] which has been used in ORBIT and many other testbeds. This framework allows easy and user-friendly management of heterogeneous hardware/software resources as well as access to APIs of different COSMOS components (e.g., the 28 GHz PAAMs and optical transport network mentioned above).

The procedure for setting up and executing an experiment in COSMOS consists of the following 5 steps (see Fig. 9):

(1) *Experiment Setup*: Reserve resources (e.g., radio and compute nodes) using the COSMOS scheduler. Use OMF commands to create and load images to the reserved nodes, and configure and bring up network links/interfaces.
(2) *Experiment Execution/Run*: Use COSMOS' control plane to send pre-prepared commands and data to the targeted nodes. Check experiment and node(s) status (e.g., success or failure) and send an "END" signal at the end of an experiment.
(3) *Data Collection*: Download or store the experimental data or live-stream the data to another node using OMF/OML commands.
(4) *Data Processing*: Process and analyze the collected data using tools such as `Python` and `MATLAB` to obtain relevant performance evaluation metrics (e.g., signal-to-noise ratio and link quality, throughput, and latency). This step can be done offline.
(5) *Resource Release*: Release the resources manually after the experiment completes or automatically after the reservation expires.

The control and management of COSMOS with 24/7 operation and timely technical support, and without causing interference to other systems, is extremely challenging. Our developed framework will support (i) efficient allocation of both the radio nodes and the heterogeneous computing resources with fast data collection and transfer methods, (ii) configuration of the Tx power and frequency bands of various radio nodes following FCC's regulations[3], and (iii) on-time monitoring of the configuration and status of the testbed when used by a large number of users worldwide.

5 DEPLOYMENT PLAN & CHALLENGES

To cope with the challenge of providing realistic environment to experimenters, COSMOS is currently being deployed in a dense urban area of about 1 sq. mile (covering about 15 city blocks and about 5 city avenues) in West Harlem, in partnership with NYC and the local community. As shown in Fig. 10, it will rely on various university and NYC assets. Particularly, Columbia University has two campuses (Morningside and Manhattanville) and CCNY has one campus, all located at the south, west, and north parts of the testbed deployment area. The center of the area includes ~20 public housing buildings and 9 public schools, managed by NYC Department of Education (DOE) and NYC Housing Authority (NYCHA), respectively. The allocation of lightpoles is managed by NYC Department of Information Technology and Telecommunications (DOITT) and NYC Department of Transportation (DOT), and requires the support of the local community board. The overall interaction with the various NYC agencies is managed by the Mayor's Office of the CTO.

---

[3] See Section 5 for the discussion regarding the FCC Innovation Zone.

Figure 10: **COSMOS' envisioned deployment area of about 1 sq. mile covering about 15 city blocks and about 5 city avenues, with tentative deployment sites marked on the map (see also Fig. 2 for the network connectivity and the three deployment phases).**

The location of the testbed was selected due to its proximity to campuses (where buildings can be used for deployment, students can be opt-in users, and university-controlled vehicles circulate), concentration of NYC assets, and the existence of an engaged community of educators, organizers, and entrepreneurs. Fig. 10 illustrates the vision for the full planned deployment with ~49 infrastructure radio nodes (9 large and ~40 medium) at the conclusion of Phase 2 (see Section 3). Large radio nodes (macro-cellular BSs) will be installed on rooftops. Medium nodes (APs or micro-cellular BSs) will be installed at street level (building side, security guard booths, and lightpoles). Finally, ~200 small (near-portable) nodes will be used as fixed or mobile devices at Columbia, CCNY, and the local community, including in vehicles (e.g., public safety cars and shuttles) and used for educational purposes. These nodes will serve as experimental network users and could be used by university students and staff, high school students, and public housing residents. Dark fiber to the various locations as indicated in Fig. 2 will be provided/deployed by Columbia, CCNY, NYC, and potentially by the PAWR industry consortium. The allocation of NYC franchisees' dark fiber is managed by the Mayor's office of the CTO and DOITT.

The deployment consists of three phases (see Fig. 2):

Pilot Phase (completed in May 2019, see Figs. 11): a proof of concept phase that targets functional verification of the designs for all three node types as well as initial performance evaluation (e.g., radio coverage planning). It included the deployment of 2 large, 3 medium, and 30 small nodes. It also included the control center at Rutgers, the central optical core switching (a Calient S320 space switch and 6 ROADM units) and computing (ToR switches and servers) at Columbia, the dark fiber to the colocation site at 32 AoA, as well as two sandboxes that provide isolated testing environments. In particular, the sandbox at Rutgers includes the Large–SB node, a medium node, and compute. It also includes a pair of Sivers IMA 60 GHz radios and a pair of InterDigital EdgeHaul 60 GHz radios. The sandbox at Columbia includes several SDRs, ToR switches, servers, and the alpha version of the IBM 28 GHz PAAMs (see Section 4.2).

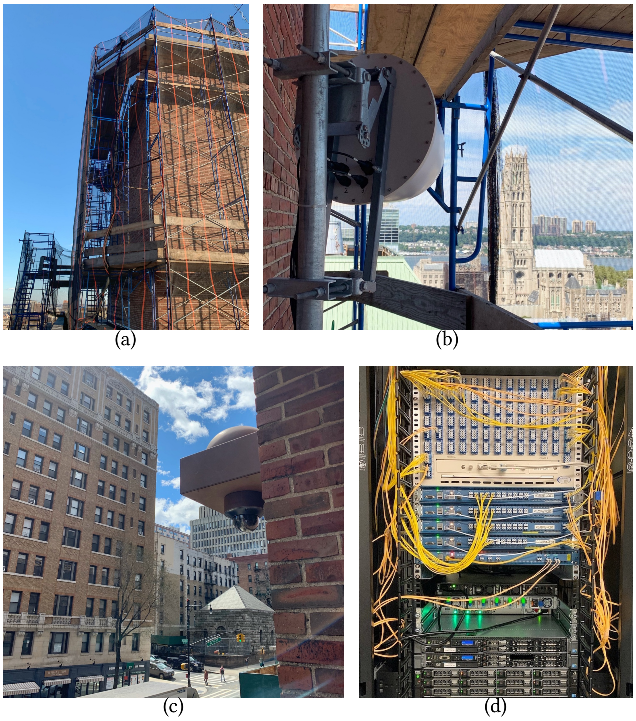
(a) (b) (c) (d)

Figure 11: **(a) The construction of the Pilot Phase large node (Large–1) on an 18th floor rooftop at Columbia, (b)–(c) Pilot Phase deployment of a large/medium node on the 18th/2nd floor (see also Fig. 4), and (d) the central optical core switching (including the Calient S320 space switch and several ROAMD units) in the data center at Columbia.**

Phase 1 (expected by the end of 2020): deployment of 3 large, 8 medium, 60 small nodes (all sub-6 GHz and with mmWave at about 20 nodes) in both Columbia and CCNY campus buildings as well as Columbia residential buildings. In addition, optical switching and computing will be deployed at the NYU data center.

Phase 2 (timeline for completion based on availability of resources): complete deployment of all nodes (all sub-6 GHz and with mmWave at ~25% of the nodes).

Throughout the deployment phases we have been conducting numerous meetings with stakeholders in NYC agencies (not limited to the ones mentioned above) and the local community (e.g., Community Boards 9 and 10, Manhattan Borough President, school districts, NYCHA tenant associations, and West Harlem Development Corporation) in order to inform them about the progress and to receive feedback about potential concerns. A technology upgrade effort is also envisioned during the testbed operation, intended to keep the COSMOS testbed current and to meet emerging user experiment requirements. The PAWR consortium is expected to provide in-kind contributions to support Phases 1 and 2 of the deployment.

Moreover, FCC has recently designated part of the COSMOS testbed area as one of the country's first two Innovation Zones [63] with allowable frequency bands of 2500-2690 MHz, 3700–4200 MHz, 5850–5925 MHz, 5925–7125 MHz, 27.5–28.35 GHz, and

Table 1: **Potential research areas that could be studied using the COSMOS testbed.**

Advanced PHY techniques: large/distributed MIMO and adaptive beamforming, full-duplex, multi-connectivity, coordinated scheduling, etc.
Conclusively verifying mmWave feasibility for mobility services
Heterogeneity, multi-homing and densification in cellular networks
Dynamic spectrum access (e.g., distributed protocols, directional sharing of mmWave frequencies)
Latency reduction in mobile network PHY/MAC and network layers
Integration of optical x-haul technology with wireless technologies
Clean slate architectures for mobile networks
Adaptive multicast for crowded venues
Edge cloud integration with wireless networks

38.6–40.0 GHz. Under this initiative, parties will have the flexibility to conduct multiple non-related experiments under a single authorization within a defined geographic area while protecting incumbent services against harmful interference. This initiative also allows experimental program license holders which are licensed to operate elsewhere to use the NYC Innovation Zone.

Overall, orchestrating the deployment plan in a realistic environment while engaging numerous university and city partners and stakeholders, and while leveraging the PAWR consortium in-kind contributions in a way that achieves the project's scientific goals poses unique challenges.

6 SUPPORTED EXPERIMENTATION

As mentioned in Sections 1 and 3, a major challenge in designing COSMOS is to enable several new classes of wireless experiments that are not supported by existing testbeds. In Table 1, we summarize the potential research areas that could be studied using the COSMOS testbed at different layers of the protocol stack (PHY to application) and in different network domains (e.g., wireless, wired, and optical). While conducting research in these areas is out of scope, prior to the community release of COSMOS, we are internally developing a representative set of example experiments that are used to drive the design and validate the capabilities and usability of the testbed. These experiments can also help other researchers envision and plan their own experiments.

We present 4 illustrative example experiments, which consider different network layers, different data paths, and different network domains (see Fig. 1). Tutorials outlining these experiments are available at [82]. We note that this is not a comprehensive list of possible experiments but rather a small subset that demonstrates and validates different capabilities.

Experiment 1: Open-Access Full-Duplex Wireless (*Data Paths 1 and 2* in Fig. 1). Full-duplex (FD) wireless – simultaneous transmission and reception at the same frequency – has the potential to double network capacity at the PHY layer and to provide many other benefits at higher layers [33–38]. Despite extensive research in this area, an open-access wireless testbed with FD-capable nodes is crucial for experimental evaluations of FD-related algorithms at the higher layers.

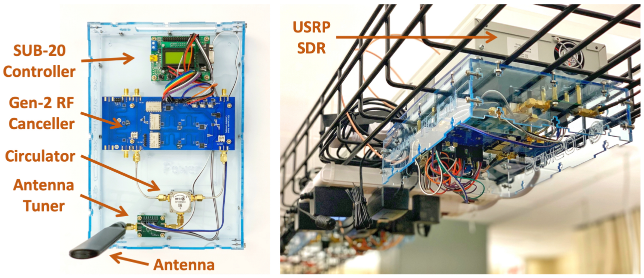

Figure 12: **The 2$^{\text{nd}}$-generation *wideband* full-duplex radio [85] (left) and its integration in the COSMOS testbed sandbox (right) allowing for remote access and compute through COSMOS' optical network [84].**

To allow the broader community to experiment with FD, following our integration of the 1$^{\text{st}}$-generation narrowband FD radio in ORBIT [83], we recently developed and integrated the 2$^{\text{nd}}$-generation *wideband* FD radios in the COSMOS sandbox 2 (SB2) [84] (see Fig. 12). This FD radio leverages the technique of frequency-domain equalization (FDE) [85] to achieve improved performance and bandwidth of RF self-interference cancellation (SIC) (e.g., 50 dB across 20 MHz). Example baseline programs and FD experiments including real-time packet-level digital SIC and packet decoding, and integration with COSMOS' optical network for remote processing using the edge cloud, are provided. Advanced example experiments at the higher layers are also under development.

We plan to integrate the 2$^{\text{nd}}$-generation wideband FD radios at various outdoor locations in COSMOS in order to support experimentation with FD in a real dense urban area. This experiment also showcases the capability of integrating customized experimental hardware in an open-access wireless testbed, where the testbed's infrastructure (e.g., SDRs and edge cloud servers) can be leveraged for different types of experiments.

Experiment 2: Optical-Wireless x-Hual Networking (*Data Paths 3 and 4* in Fig. 1). An optical-wireless x-haul experiment has been developed to demonstrate the C-RAN architecture and its integration with the optical x-haul network and SDN control. The experimental setup is shown in Fig. 6(b). In particular, the wideband FD radio, described above, serves as an FD BS located in SB2. The BS sends 20 MHz baseband I/Q data (limited by the interface between the SDR and server) over the dark fiber (with 10/100 Gbps transceivers) to the optical switches at the remote data center of NYU at 32 AoA. It is then sent back to the data center at Columbia through the red route for remote digital signal processing (∼14 miles, Data Path 4).

Simultaneously, a video multicast application [86] operates on the x-haul network on a different optical wavelength. The AP, which needs to dynamically adapt to the channel conditions of several users, receives video streams from two servers through *on-demand optical switching* managed by a customized Ryu-based SDN controller [78, 87]. In particular, one is received from an edge cloud server through the short green route (∼0.3 mile, Data Path 3), and the other from a central cloud server through the long red route (∼14 miles, Data Path 4). Due to its high capacity (10s of Gbps) and low latency (<1 ms), such an optical x-haul network can support high-performance computing tasks and flexible topologies. The experiment evaluates the capability of the multicast application to switch between local and remote servers (via optical support) while responding to the users' channel states.

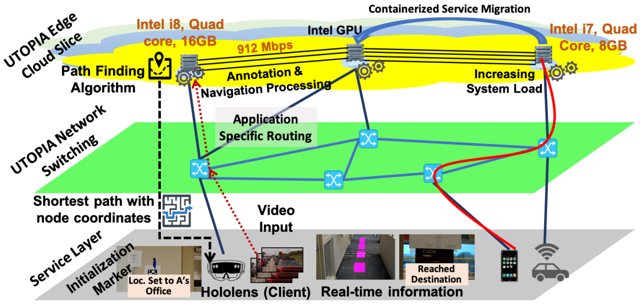

Figure 13: **Setup of the low-latency AR edge cloud experiment, where the mesh topology can be configured by COSMOS' optical network.**

Table 2: **Comparison of the AR application end-to-end latency.**

Normalized Average System Load	Same Edge Cloud	Edge Cloud w/ CSM	Edge Cloud w/ CSM and ASR
0.1	21.6 ms	21.6 ms	21.2 ms
0.4	89.7 ms	75.7 ms	58.5 ms
0.7	148.5 ms	110.0 ms	92.1 ms

These optical-wireless x-haul networking experiments are examples of a broader class of experiments that can be uniquely supported by COSMOS' programmable optical network, which can be dynamically configured by an SDN controller to emulate different network topologies and functionalities (e.g., C-RAN). Additional examples with details are available in [88, 89].

Experiment 3: Low-Latency AR Edge Cloud (*Data Path 3* in Fig. 1). Edge cloud promises to support low-latency applications by bringing computing close to the user. As a distributed computing infrastructure, it poses challenges in managing heterogeneous bandwidth and computing resources, and providing seamless service mobility. Future AR/VR applications will heavily rely on edge cloud to fulfill their intensive computing requirements in real-time, which requires techniques such as application specific routing (ASR) [90].

In order to evaluate the feasibility of AR-based low-latency annotation and smart navigation applications, we developed an example experiment in COSMOS. Fig. 13 shows the experimental setup, where different system loads, inter-edge bandwidth, and containerized services for migrations (CSM) [91, 92] are configured using COSMOS' heterogeneous computing resources. For a mix of AR applications, essential edge cloud evaluation metrics (e.g., application delay) are obtained. Table 2 shows the end-to-end latency comparison with and without using the CSM and ASR techniques, where the latency depends on both the computation capability and inter-edge bandwidth. We also plan to evaluate fast service migrations in more complex network topologies using COSMOS' reconfigurable optical network.

This is an example experiment where various requirements posed by the AR applications can be fulfilled using a large-scale testbed with a tightly coupled edge cloud for providing low network delay, reduced processing overhead, and inter-edge cloud optical connectivity enabling techniques (e.g., containerized service migration to handle dynamic system load and mobility). In addition, COSMOS allows for experimental evaluation of the end-to-end service performance in an outdoor dense urban environment, which is currently not supported by other testbeds.

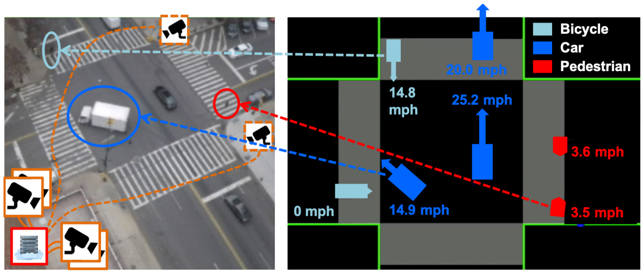

(a) (b)

Figure 14: **Locations of the cameras and edge cloud servers at the COSMOS pilot intersection with an example "radar-screen" captured by a camera deployed on the 12th floor, and the corresponding locations and velocity vectors of bicycles/cars/pedestrians in the "radar-screen" obtained using deep learning algorithms.**

Experiment 4: Smart City Intersections (*Data Path 2* in Fig. 1). The deployment of autonomous vehicles in dense urban environments presents unique challenges such as vehicles moving at various speeds, obstructions that are opaque to in-vehicle sensors, and erratic pedestrians. Smart city intersection will be at the core of an AI-powered traffic management system for metropolises. COSMOS will provide all components needed for developing smart intersections and for supporting cloud-connected vehicles. In particular, COSMOS will enable vehicles to wirelessly share in-vehicle sensor data with other vehicles and the edge cloud servers. Data collected by infrastructure sensors deployed in COSMOS (e.g., street-level and bird's eye cameras) will be aggregated at the edge cloud servers, which can run real-time algorithms to detect and track all objects for traffic monitoring and management.

We developed an example experiment using the cameras and edge cloud servers deployed during the Pilot Phase [93]. In particular, we devised customized deep learning algorithms to create a "radar-screen" movie tracking all objects in the intersection. The algorithms are capable of detecting objects of notably different sizes observed from the bird's eye view. Fig. 14 shows an example snapshot of the results, where the locations and velocity vectors of bicycles/cars/pedestrians are indicated in different colors. We plan to dynamically distribute the video analytics to various COSMOS' compute nodes based on bandwidth and latency requirements. The images and video streams collected by these cameras, accompanied with COSMOS' edge cloud and heterogeneous computing resources, can be used for various sensing and monitoring applications spanning smart cities and intersections. This experiment is the first step towards other experiments that will consider coordination between several intersections both for network control and traffic control.

7 OUTREACH & EDUCATION

A major challenge in building an urban testbed is to provide benefits to not only the research community but also the local community, where K–12 students come from diverse racial, ethnic, and socioeconomic backgrounds. The curricula in such an area must take into consideration the urban space in which the students reside as well as connect it to science and engineering practices that use the tenets of inquiry-based teaching and learning.

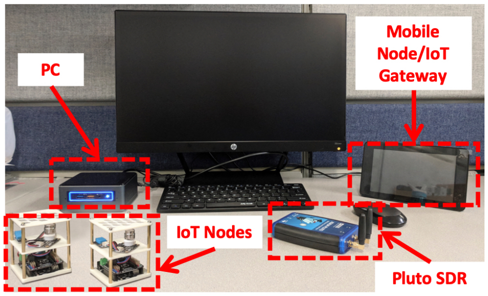

Figure 15: **The COSMOS Education Toolkit which has been used in 18 schools to run educational lab experiments.**

To cope with this challenge, we focus on transforming COSMOS into an innovative learning platform for middle and high school students that would help bridge the digital divide and provide significant educational benefits for the local community. Hence, COSMOS' education plan includes enhancing STEM teachers' professional development and collaborating with them to create a set of hands-on COSMOS-based educational lab experiments to enhance STEM education. Specifically, we organized a professional development program for NYC teachers which consists of three phases: (i) a two-week structured learning phase, during which ten teachers attend lectures about wireless communications and networking, and perform various hands-on experiments using SDRs and IoT nodes, (ii) a four-week research phase, during which the teachers collaborate with COSMOS faculty, postdocs, and graduate students in order to co-create educational lab experiments aligned with the K–12 Next Generation Science Standards (NGSS), and (iii) support for the teachers/students over the academic year for executing the wireless-based experiments.

The outcomes of the program can be summarized as follows: (i) the teachers gained in-depth knowledge of wireless communications and networking, (ii) the development (by the teachers and the project team) of over 100 NGSS-aligned labs (in the areas of math, science, and computer science) that will be used throughout the academic year, and (iii) the development of the COSMOS Education Toolkit (see see Fig. 15) which is a small pre-configured COSMOS node that will provide the necessary infrastructure for the students to execute the lab experiments locally or by connecting to COSMOS (for more details see [64]). The software executes and manages the experiments using an SDR and IoT devices in the same operational philosophy as COSMOS, adding easy-to-use enhancements, as it is designed for non-technical school teachers and students.

The professional development program has been implemented during the the summers of 2018 and 2019, with 20 participants who used the COSMOS Education Toolkit in 18 middle and high schools throughout NYC. Similar programs will be carried out in the future with the goal of allowing any public school in the city to remotely use the testbed for educational experiments.

8 CONCLUSIONS

This paper outlines the challenges and partial experience gained in designing a 1 sq. mile city-scale *programmable* advanced wireless testbed, COSMOS, which incorporates mmWave radios, optical transport, and edge cloud. It outlines our approaches to addressing many of the challenges and provides insight into the research areas that can be studied in the testbed once operational. It also discusses the unique challenges of deploying the testbed in a dense urban environment and the opportunity to use is as an outreach and educational tool. In summary, the COSMOS testbed will offer the following unique capabilities:

1. A large number of state-of-the-art SDRs with three form factors (large, medium, and small) for different deployment scenarios (rooftop, small cell/lightpole, and user equipment), which are equipped with sub-6 GHz and mmWave front ends.
2. A unique first-of-its-kind programmable high-bandwidth optical transport network using long-range multi-hop dark fiber (with distances of 10s of miles) to emulate various optical front/mid/back-haul (x-haul) networks.
3. Core and edge cloud equipped with programmable heterogeneous computing resources (CPUs, GPUs, and FPGAs), which is tightly coupled with radio nodes and optical transport network.
4. Support for a wide range of research and experiments at-scale within an FCC Innovation Zone, which include (but not limited to): (i) integration of customized experimental hardware, (ii) experimentation with emerging Physical layer technologies (e.g., full-duplex wireless, mmWave with multi-beam support and agile beam-steering capabilities, distributed and adaptive beamforming), and (iii) experimentation with algorithms and applications at the higher layers (e.g., edge computing, AR/VR, and smart cities and intersections).

ACKNOWLEDGMENTS

COSMOS is part of the NSF's PAWR program and is funded in part by NSF award CNS-1827923 and by the PAWR Industry Consortium. The experiments in Section 6 were supported in part by NSF Grants CNS-1345295, CNS-1650685, CNS-1650669, ECCS-1547406, and by AT&T. The education program described in Section 7 was supported in part by supplements to NSF awards EFMA-1641100 (NewLAW EFRI), CNS-1329939, CNS-1513110, CNS-1527750, CNS-1650669, CNS-1650685, CNS-1702952, CNS-1730043, ECCS-1547332, ECCS-1547406, OAC-1541069, Columbia Data Science Institute, and by AT&T foundation.

We thank the following for their numerous important contributions to the COSMOS testbed: Janice Campanella, Prasanthi Maddala, Nilanjan Paul, Newman Wilson (WINLAB, Rutgers); Henning Schulzrinne (CS, Columbia); Sharon Sputz (Data Science Institute, Columbia); Tianwei Deng, Manav Kohli, Jonathan Ostrometzky (EE, Columbia); Karen Cheng, Emily Ford (Engineering Outreach, Columbia); Anthony Avendano (Facilities and Operation, Columbia); Flores Forbes, Victoria Mason-Ailey (Government and Community Affairs, Columbia); Alan Crosswell, Daniel Gaitings, Thomas Rom (CUIT); Sheila Borges, Ben Esner (NYU Center for K12 STEM Education); Shivendra Panwar, Sundeep Rangan (ECE, NYU); Joshua Breitbart (New York City Mayor's Office of the CTO); Clayton Banks, Bruce Lincoln (Silicon Harlem); Jiakai Yu, Shengxiang Zhu (College of Optical Sciences, U. Arizona); Myung Lee, Rosemarie Wesson (CCNY); Arun Paidimarri, Bodhisatwa Sadhu, Alberto Valdes-Garcia (IBM Research); Nikos Makris (U. Thessaly); and Artur Minakhmetov (Télécom Paris). We thank several other contributors to the testbed in Rutgers, Columbia, NYU, NYC, Silicon Harlem, U. Arizona, CCNY, and IBM Research.

We thank the anonymous reviewers and shepherd for their insightful comments.

REFERENCES

[1] Cisco. The Zettabyte era: Trends and analysis. https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html, 2018.

[2] Jeffrey G Andrews, Stefano Buzzi, Wan Choi, Stephen V Hanly, Angel Lozano, Anthony CK Soong, and Jianzhong Charlie Zhang. What will 5G be? *IEEE J. Sel. Areas Commun.*, 32(6):1065–1082, 2014.

[3] Cheng-Xiang Wang, Fourat Haider, Xiqi Gao, Xiao-Hu You, Yang Yang, Dongfeng Yuan, Hadi M Aggoune, Harald Haas, Simon Fletcher, and Erol Hepsaydir. Cellular architecture and key technologies for 5G wireless communication networks. *IEEE Commun. Mag.*, 52(2):122–130, 2014.

[4] Dipankar Raychaudhuri and Narayan B Mandayam. Frontiers of wireless and mobile communications. *Proc. IEEE*, 100(4):824–840, 2012.

[5] NGMN-Alliance. 5G white paper. *Next generation mobile networks, white paper*, 2015.

[6] Ericsson. 5G radio access. https://www.ericsson.com/assets/local/publications/white-papers/wp-5g.pdf, 2016.

[7] Qualcomm. 5G: The fabric for society. https://www.qualcomm.com/media/documents/files/5g-vision-use-cases.pdf, 2018.

[8] Mansoor Shafi, Andreas F Molisch, Peter J Smith, Thomas Haustein, Peiying Zhu, Prasan De Silva, Fredrik Tufvesson, Anass Benjebbour, and Gerhard Wunder. 5G: A tutorial overview of standards, trials, challenges, deployment, and practice. *IEEE J. Sel. Areas Commun.*, 35(6):1201–1221, 2017.

[9] Mamta Agiwal, Abhishek Roy, and Navrati Saxena. Next generation 5G wireless networks: A comprehensive survey. *Commun. Surveys Tuts.*, 18(3):1617–1655, 2016.

[10] Kevin Boos, David Chu, and Eduardo Cuervo. Flashback: Immersive virtual reality on mobile devices via rendering memoization. In *Proc. ACM MobiSys'16*, 2016.

[11] Luyang Liu, Hongyu Li, and Marco Gruteser. Edge assisted real-time object detection for mobile augmented reality. In *Proc. ACM MobiCom'19*, 2019.

[12] Eduardo Cuervo, Krishna Chintalapudi, and Manikanta Kotaru. Creating the perfect illusion: What will it take to create life-like virtual reality headsets? In *Proc. ACM HotMobile'18*, 2018.

[13] Luyang Liu, Ruiguang Zhong, Wuyang Zhang, Yunxin Liu, Jiansong Zhang, Lintao Zhang, and Marco Gruteser. Cutting the cord: Designing a high-quality untethered VR system with low latency remote rendering. In *Proc. ACM MobiSys'18*, 2018.

[14] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. Enabling high-quality untethered virtual reality. In *Proc. USENIX NSDI'17*, 2017.

[15] Sumit Maheshwari, Dipankar Raychaudhuri, Ivan Seskar, and Francesco Bronzino. Scalability and performance evaluation of edge cloud systems for latency constrained applications. In *Proc. IEEE/ACM Symposium on Edge Computing (SEC)*, 2018.

[16] Md Whaiduzzaman, Mehdi Sookhak, Abdullah Gani, and Rajkumar Buyya. A survey on vehicular cloud computing. *J. Network and Computer Applications*, 40: 325–344, 2014.

[17] Kengo Sasaki, Naoya Suzuki, Satoshi Makido, and Akihiro Nakao. Vehicle control system coordinated between cloud and mobile edge computing. In *Proc. IEEE SICE'16*, 2016.

[18] Platforms for advanced wireless research (PAWR). https://www.advancedwireless.org/, 2019.

[19] PAWR: Why now? https://www.advancedwireless.org/why-now/, 2019.

[20] Cloud enhanced open software defined mobile wireless testbed for city-scale deployment (COSMOS). https://cosmos-lab.org/, 2019.

[21] Liangping Ma, Xiaofeng Han, and Chien-Chung Shen. Dynamic open spectrum sharing MAC protocol for wireless ad hoc networks. In *Proc. IEEE DySPAN'05*, 2005.

[22] Zhu Ji and KJ Ray Liu. Cognitive radios for dynamic spectrum access-dynamic spectrum sharing: A game theoretical overview. *IEEE Commun. Mag.*, 45(5):88–94, 2007.

[23] Sudeep Bhattarai, Jung-Min Jerry Park, Bo Gao, Kaigui Bian, and William Lehr. An overview of dynamic spectrum sharing: Ongoing initiatives, challenges, and a roadmap for future research. *IEEE Trans. Cognitive Commun. Netw.*, 2(2):110–128, 2016.

[24] Theodore S Rappaport, Shu Sun, Rimma Mayzus, Hang Zhao, Yaniv Azar, Kevin Wang, George N Wong, Jocelyn K Schulz, Mathew Samimi, and Felix Gutierrez. Millimeter-wave mobile communications for 5G cellular: It will work! *IEEE Access*, 1:335–349, 2013.

[25] Sundeep Rangan, Theodore S Rappaport, and Elza Erkip. Millimeter-wave cellular wireless networks: Potentials and challenges. *Proc. IEEE*, 102(3):366–385, 2014.

[26] Joan Palacios, Danilo De Donno, Domenico Giustiniano, and Joerg Widmer. Speeding up mmWave beam training through low-complexity hybrid transceivers. In *Proc. IEEE PIMRC'16*, 2016.

[27] Omid Abari, Haitham Hassanieh, Michael Rodriguez, and Dina Katabi. Millimeter wave communications: From point-to-point links to agile network connections. In *Proc. ACM HotNets'16*, 2016.

[28] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. BeamSpy: Enabling robust 60 GHz links under blockage. In *Proc. USENIX NSDI'16*, 2016.

[29] Teng Wei and Xinyu Zhang. Pose information assisted 60 GHz networks: Towards seamless coverage and mobility support. In *Proc. ACM MobiCom'17*, 2017.

[30] Margarita Gapeyenko, Andrey Samuylov, Mikhail Gerasimenko, Dmitri Moltchanov, Sarabjot Singh, Mustafa Riza Akdeniz, Ehsan Aryafar, Nageen Himayat, Sergey Andreev, and Yevgeni Koucheryavy. On the temporal effects of mobile blockers in urban millimeter-wave cellular scenarios. *IEEE Trans. Veh. Technol.*, 66(11):10124–10138, 2017.

[31] Muhammad Kumail Haider and Edward W Knightly. Mobility resilience and overhead constrained adaptation in directional 60 GHz WLANs: Protocol design and system implementation. In *Proc. ACM MobiHoc'16*, 2016.

[32] Yasaman Ghasempour, Claudio RCM da Silva, Carlos Cordeiro, and Edward W Knightly. IEEE 802.11 ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi. *IEEE Commun. Mag.*, 55(12):186–192, 2017.

[33] Jung Il Choi, Mayank Jain, Kannan Srinivasan, Phil Levis, and Sachin Katti. Achieving single channel, full duplex wireless communication. In *Proc. ACM MobiCom'10*, 2010.

[34] Bozidar Radunovic, Dinan Gunawardena, Peter Key, Alexandre Proutiere, Nikhil Singh, Vlad Balan, and Gerald Dejean. Rethinking indoor wireless mesh design: Low power, low frequency, full-duplex. In *Proc. IEEE WiMesh'10*, 2010.

[35] Dinesh Bharadia, Emily McMilin, and Sachin Katti. Full duplex radios. In *Proc. ACM SIGCOMM'13*, 2013.

[36] Ashutosh Sabharwal, Philip Schniter, Dongning Guo, Daniel W Bliss, Sampath Rangarajan, and Risto Wichman. In-band full-duplex wireless: Challenges and opportunities. *IEEE J. Sel. Areas Commun.*, 32(9):1637–1652, 2014.

[37] Harish Krishnaswamy and Gil Zussman. 1 chip 2x the bandwidth. *IEEE Spectrum*, 53(7):38–54, 2016.

[38] Jin Zhou, Negar Reiskarimian, Jelena Diakonikolas, Tolga Dinc, Tingjun Chen, Gil Zussman, and Harish Krishnaswamy. Integrated full duplex radios. *IEEE Commun. Mag.*, 55(4):142–151, 2017.

[39] Erik G Larsson, Ove Edfors, Fredrik Tufvesson, and Thomas L Marzetta. Massive MIMO for next generation wireless systems. *IEEE Commun. Mag.*, 52(2):186–195, 2014.

[40] Xing Zhang, John Tadrous, Evan Everett, Feng Xue, and Ashutosh Sabharwal. Angle-of-arrival based beamforming for FDD massive MIMO. In *Proc. Asilomar'15*, 2015.

[41] Kate Ching-Ju Lin, Shyamnath Gollakota, and Dina Katabi. Random access heterogeneous MIMO networks. In *Proc. ACM SIGCOMM'11*, 2011.

[42] Lu Lu, Geoffrey Ye Li, A Lee Swindlehurst, Alexei Ashikhmin, and Rui Zhang. An overview of massive MIMO: Benefits and challenges. *IEEE J. Sel. Topics Signal Process.*, 8(5):742–758, 2014.

[43] Ehsan Aryafar, Mohammad Amir Khojastepour, Karthikeyan Sundaresan, Sampath Rangarajan, and Mung Chiang. MIDU: Enabling MIMO full duplex. In *Proc. ACM MobiCom'12*, 2012.

[44] Swarun Kumar, Diego Cifuentes, Shyamnath Gollakota, and Dina Katabi. Bringing cross-layer MIMO to today's wireless LANs. In *Proc. ACM SIGCOMM'13*, 2013.

[45] Xiufeng Xie, Eugene Chai, Xinyu Zhang, Karthikeyan Sundaresan, Amir Khojastepour, and Sampath Rangarajan. Hekaton: Efficient and practical large-scale MIMO. In *Proc. ACM MobiCom'15*, 2015.

[46] Wei-Liang Shen, Kate Ching-Ju Lin, Ming-Syan Chen, and Kun Tan. SIEVE: Scalable user grouping for large MU-MIMO systems. In *Proc. IEEE INFOCOM'15*, 2015.

[47] Clayton Shepard, Hang Yu, Narendra Anand, Erran Li, Thomas Marzetta, Richard Yang, and Lin Zhong. Argos: Practical many-antenna base stations. In *Proc. ACM MobiCom'12*, 2012.

[48] Clayton Shepard, Hang Yu, and Lin Zhong. ArgosV2: A flexible many-antenna research platform. In *Proc. ACM MobiCom'13*, 2013.

[49] Christopher Husmann, Georgios Georgis, Konstantinos Nikitopoulos, and Kyle Jamieson. FlexCore: Massively parallel and flexible processing for large MIMO access points. In *Proc. USENIX NSDI'17*, 2017.

[50] Yuanjie Li, Zengwen Yuan, and Chunyi Peng. A control-plane perspective on reducing data access latency in LTE networks. In *Proc. ACM MobiCom'17*, 2017.

[51] Haoyang Wu, Tao Wang, Zengwen Yuan, Chunyi Peng, Zhiwei Li, Zhaowei Tan, Boyan Ding, Xiaoguang Li, Yuanjie Li, Jun Liu, et al. The tick programmable low-latency SDR system. In *Proc. ACM MobiCom'17*, 2017.

[52] Mahadev Satyanarayanan, Zhuo Chen, Kiryong Ha, Wenlu Hu, Wolfgang Richter, and Padmanabhan Pillai. Cloudlets: At the leading edge of mobile-cloud convergence. In *Proc. MobiCASE'14*, 2014.

[53] Tan Zhang, Aakanksha Chowdhery, Paramvir Victor Bahl, Kyle Jamieson, and Suman Banerjee. The design and implementation of a wireless video surveillance system. In *Proc. ACM MobiCom'15*, 2015.

[54] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.

[55] Ganesh Ananthanarayanan, Paramvir Bahl, Peter Bodík, Krishna Chintalapudi, Matthai Philipose, Lenin Ravindranath, and Sudipta Sinha. Real-time video analytics: The killer app for edge computing. *IEEE Computer*, 50(10):58–67, 2017.

[56] Swati Rallapalli, Aishwarya Ganesan, Krishna Chintalapudi, Venkat N Padmanabhan, and Lili Qiu. Enabling physical analytics in retail stores using smart glasses. In *Proc. ACM MobiCom'14*, 2014.

[57] Chien-Chun Hung, Ganesh Ananthanarayanan, Peter Bodik, Leana Golubchik, Minlan Yu, Paramvir Bahl, and Matthai Philipose. VideoEdge: Processing camera streams using hierarchical clusters. In *Proc. IEEE/ACM SEC'18*, 2018.

[58] Weisong Shi, Jie Cao, Quan Zhang, Youhuizi Li, and Lanyu Xu. Edge computing: Vision and challenges. *IEEE Internet Things J.*, 3(5):637–646, 2016.

[59] Peng Liu, Dale Willis, and Suman Banerjee. Paradrop: Enabling lightweight multi-tenancy at the network's extreme edge. In *Proc. IEEE/ACM SEC'16*, 2016.

[60] Leandro Y Mano, Bruno S Faiçal, Luis HV Nakamura, Pedro H Gomes, Giampaolo L Libralon, Rodolfo I Meneguete, PR Geraldo Filho, Gabriel T Giancristofaro, Gustavo Pessin, Bhaskar Krishnamachari, et al. Exploiting IoT technologies for enhancing health smart homes through patient identification and emotion recognition. *Computer Communications*, 89:178–190, 2016.

[61] Agusti Solanas, Constantinos Patsakis, Mauro Conti, Ioannis S Vlachos, Victoria Ramos, Francisco Falcone, Octavian Postolache, Pablo A Pérez-Martínez, Roberto Di Pietro, Despina N Perrea, et al. Smart health: A context-aware health paradigm within smart cities. *IEEE Commun. Mag.*, 52(8):74–81, 2014.

[62] Carolina Tripp Barba, Miguel Angel Mateos, Pablo Reganas Soto, Ahmad Mohamad Mezher, and Mónica Aguilar Igartua. Smart city for VANETs using warning messages, traffic statistics and intelligent traffic lights. In *Proc. IEEE Intelligent Vehicles Symposium*, 2012.

[63] Innovation zones for program experimental licenses in designated portions of new york city and salt lake city. https://docs.fcc.gov/public/attachments/DA-19-923A1.pdf/, 2019.

[64] COSMOS education toolkit. https://cosmos-lab.org/cosmos-toolkit/, 2019.

[65] Bristol Is Open. http://www.bristolisopen.com/.

[66] ADRENALINE testbed. http://networks.cttc.es/ons/adrenaline/sdnnfv-cloud-computing-platform-and-core-network-for-5g-services/.

[67] POWDER-RENEW. https://powderwireless.net/.

[68] Robert Ricci, Eric Eide, and CloudLab Team. Introducing CloudLab: Scientific infrastructure for advancing cloud architectures and applications. *The USENIX Magazine*, 39(6):36–38, 2014.

[69] Mark Berman, Jeffrey S Chase, Lawrence Landweber, Akihiro Nakao, Max Ott, Dipankar Raychaudhuri, Robert Ricci, and Ivan Seskar. GENI: A federated testbed for innovative network experiments. *Computer Networks*, 61:5–23, 2014.

[70] Eric Eide, Leigh Stoller, and Jay Lepreau. An experimentation workbench for replayable networking research. In *Proc. USENIX NSDI'07*, 2007.

[71] Brent Chun, David Culler, Timothy Roscoe, Andy Bavier, Larry Peterson, Mike Wawrzoniak, and Mic Bowman. PlanetLab: An overlay testbed for broad-coverage services. *ACM SIGCOMM Computer Communication Review*, 33(3):3–12, 2003.

[72] Serge Fdida, Timur Friedman, and Thierry Parmentelat. OneLab: An open federated facility for experimentally driven future internet research. In *New Network Architectures*, pages 141–152. Springer, 2010.

[73] The CIAN TOAN testbed. http://cian-erc.uawebhost.arizona.edu/testbed-capabilities.

[74] Dipankar Raychaudhuri, Ivan Seskar, Max Ott, Sachin Ganu, Kishore Ramachandran, Haris Kremo, Robert Siracusa, Hang Liu, and Manpreet Singh. Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols. In *Proc. IEEE WCNC'05*, 2005.

[75] Bodhisatwa Sadhu, Yahya Tousi, Joakim Hallin, Stefan Sahl, Scott K Reynolds, Örjan Renström, Kristoffer Sjögren, Olov Haapalahti, Nadav Mazor, Bo Bokinge, et al. A 28-GHz 32-element TRX phased-array IC with concurrent dual-polarized operation and orthogonal phase and gain control for 5G communications. *IEEE J. Solid-State Circuits*, 52(12):3373–3391, 2017.

[76] Tingjun Chen, Manav Kohli, Tianyi Dai, Angel Daniel Estigarribia, Dmitry Chizhik, Jinfeng Du, Rodolfo Feick, Reinaldo A Valenzuela, and Gil Zussman. 28 GHz channel measurements in the COSMOS testbed deployment area. In *Proc. ACM MobiCom'19 Workshop on Millimeter-Wave Networks and Sensing System (mmNets)*, 2019.

[77] Jiakai Yu, Tingjun Chen, Craig Gutterman, Shengxiang Zhu, Gil Zussman, Ivan Seskar, and Dan Kilper. COSMOS: Optical architecture and prototyping. In *Proc. OSA OFC'19*, 2019.

[78] Jiakai Yu, Craig Gutterman, Artur Minakhmetov, Michael Sherman, Tingjun Chen, Shengxiang Zhu, Gil Zussman, Ivan Seskar, and Dan Kilper. Dual use SDN controller for management and experimentation in a field deployed testbed. In *Proc. OSA OFC'20, T3J.3*, 2020.

[79] Andy Bavier, Nick Feamster, Mark Huang, Larry Peterson, and Jennifer Rexford. In VINI veritas: Realistic and controlled network experimentation. In *Proc. ACM SIGCOMM'06*, 2006.

[80] Raul Munoz, Ricard Vilalta, Ramon Casellas, Ricardo Martínez, Thomas Szyrkowiec, Achim Autenrieth, Victor López, and Diego López. SDN/NFV orchestration for dynamic deployment of virtual SDN controllers as VNF for multi-tenant optical networks. In *Proc. OSA OFC'15*, 2015.

[81] Thierry Rakotoarivelo, Maximilian Ott, Guillaume Jourjon, and Ivan Seskar. OMF: A control and management framework for networking testbeds. *ACM SIGOPS Operating Systems Review*, 43(4):54–59, 2010.

[82] COSMOS tutorials. https://wiki.cosmos-lab.org/wiki/tutorials/, 2019.

[83] Tingjun Chen, Mahmood Baraani Dastjerdi, Guy Farkash, Jin Zhou, Harish Krishnaswamy, and Gil Zussman. Open-access full-duplex wireless in the ORBIT testbed. *arXiv preprint arXiv:1801.03069v2*, 2018.

[84] Craig Gutterman, Artur Minakhmetov, Jiakai Yu, Michael Sherman, Tingjun Chen, Shengxiang Zhu, Ivan Seskar, Dipankar Raychaudhuri, Dan Kilper, and Gil Zussman. Experimentation with full-duplex wireless in the COSMOS testbed. In *Proc. IEEE ICNP'19 Workshop on Midscale Education and Research Infrastructure and Tools (MERIT)*, 2019.

[85] Tingjun Chen, Mahmood Baraani Dastjerdi, Jin Zhou, Harish Krishnaswamy, and Gil Zussman. Wideband full-duplex wireless via frequency-domain equalization: Design and experimentation. In *Proc. ACM MobiCom'19*, 2019.

[86] Varun Gupta, Craig Gutterman, Yigal Bejerano, and Gil Zussman. Experimental evaluation of large scale WiFi multicast rate control. *IEEE Trans. Wireless Commun.*, 17(4):2319–2332, 2018.

[87] Yao Li, Weiyang Mo, Shengxiang Zhu, Yiwen Shen, Jiakai Yu, Payman Samadi, Keren Bergman, and Daniel C Kilper. tSDX: Enabling impairment-aware all-optical inter-domain exchange. *J. Lightwave Technol.*, 36(1):142–154, 2018.

[88] Artur Minakhmetov, Craig Gutterman, Tingjun Chen, Cedric Ware, Luigi Iannone, Dan Kilper, and Gil Zussman. Experiments on cloud-RAN wireless handover using optical switching in a dense urban testbed. In *Proc. OSA OFC'20, Th2A.25*, 2020.

[89] Tingjun Chen, Jackson Welles, Manav Kohli, Mahmood Baraani Dastjerdi, Jakub Kolodziejski, Michael Sherman, Ivan Seskar, Harish Krishnaswamy, and Gil Zussman. Programmable optical x-haul network in the COSMOS testbed. In *Proc. IEEE ICNP'19 Workshop on Midscale Education and Research Infrastructure and Tools (MERIT)*, 2019.

[90] Francesco Bronzino, Sumit Maheshwari, Ivan Seskar, and Dipankar Raychaudhuri. Novn: named-object based virtual network architecture. In *Proc. ACM ICDCN'19*, 2019.

[91] Sumit Maheshwari, Shalini Choudhury, Ivan Seskar, and Dipankar Raychaudhuri. Traffic-aware dynamic container migration for real-time support in mobile edge clouds. In *Proc. IEEE ANTS'18*, 2018.

[92] Sumit Maheshwari, Wuyang Zhang, Ivan Seskar, Yanyong Zhang, and Dipankar Raychaudhuri. EdgeDrive: Supporting advanced driver assistance systems using mobile edge clouds networks. In *Proc. IEEE INFOCOM'19 Workshop on Big Data and Cloud Performance (DCPerf)*, 2019.

[93] Shiyun Yang, Emily Bailey, Zhengye Yang, Jonatan Ostrometzky, Gil Zussman, Ivan Seskar, and Zoran Kostic. COSMOS smart intersection: Edge compute and communications for bird's eye object tracking. In *Proc. 4th International Workshop on Smart Edge Computing and Networking (SmartEdge'20)*, 2020.