A Software-Defined Programmable Testbed for Beyond 5G Optical-Wireless Experimentation at City-Scale

Tingjun Chen, Jiakai Yu, Arthur Minakhmetov, Craig Gutterman, Michael Sherman, Shengxiang Zhu, Steven Santaniello, Aishik Biswas, Ivan Seskar, Gil Zussman, and Dan Kilper

ABSTRACT

Wireless traffic will significantly increase in next-generation networks. Therefore, there is a need for novel optical network front-/mid-/backhaul (x-haul) architectures that can deliver ultrahigh bandwidth and ultra-low-latency services to wireless and edge cloud systems, as well as support various cloud radio access network (C-RAN) architectures. Designing and evaluating such architectures requires large-scale experimental platforms. In this article, we present the design and implementation of a unique, remotely accessible, disaggregated, and programmable optical network as the key component of the city-scale PAWR COSMOS advanced wireless testbed. We discuss the design and implementation of a dedicated multi-functional Ryu software-defined networking (SDN) controller for the testbed's optical network wavelength channel assignment and topology reconfiguration. We present an example use of the SDN controller for monitoring and automated measurement of the wavelength-dependent amplifier gain in deployed reconfigurable optical add/drop multiplexer (ROADM) units. We also demonstrate a pilot experiment focusing on a wireless handover scheme for C-RAN architectures in the COSMOS testbed, where high traffic volume can be supported by dynamic wavelength allocation via optical switching in the fronthaul. The field-deployed testbed and SDN controller can be used by the research community to push the envelope in the area of optical networks for beyond 5G and 6G.

Introduction

The design of beyond 5G and 6G mobile networks is driven by rapidly growing wireless traffic stemming from diverse services and applications, such as the Internet of Things, connected vehicles, and virtual reality. Many of these applications require ultra-high bandwidth and ultra-low latency not only in the wireless access, but also throughout the optical front-/mid-/backhaul (x-haul) network connecting to the (edge) cloud. This will result in significantly increased bandwidth requirements from the underlying optical x-haul networks and motivates tight integration of the wireless, optical, and computing infrastructure [1]. Specifically, core and

metro networks, scalable to large size and capacity, rely on amplified reconfigurable optical add/drop multiplexer (ROADM) transmission systems, which suffer from complex engineering and slow (i.e., minutes) reconfiguration times. Point-to-point and passive optical network (PON) systems used in access networks are highly adaptive, but limited in size and capacity. Thus, there is a growing need to extend the capabilities of these existing systems and develop new technologies to better address the challenges posed by emerging wireless networks.

Figure 1 shows an example of a cloud radio access network (C-RAN) architecture, where substantial baseband processing tasks can be offloaded to powerful edge cloud or computing clusters via the fronthaul. In particular, distributed units (DUs) and centralized units (CUs) in 5G C-RANs will be placed in centralized locations (e.g., the edge cloud) to allow sharing of computing resources among densely deployed radio units (RUs). In addition, radio over fiber (RoF) between the RU and DU/CU will require significant optical bandwidth (e.g., a 10 Gb/s wireless link will require over 150 Gb/s capacity in the optical fronthaul [2]). This calls for new types of optical networks, such as time slotted optical switching networks, for efficient resource allocation. As illustrated in Fig. 1, optical space switching supporting fiber and wavelength switching are two key switching modalities under investigation in a wide range of architectures [3]. These technologies allow for novel multi-dimensional optical transport systems, in which wavelengths with high capacities can be provisioned and dynamically allocated to satisfy the different bandwidth and computing requirements of the RUs. In addition to evaluating the performance of different optical architecture, studying the trade-offs related to integration of wireless and optical systems and to data processing locations is essential.

As optical networks are used in new applications such as 5G C-RAN (Fig. 1) and built out into larger, more complex topologies, the number of wavelengths and switch dimensions will continue to increase, placing greater demands on the control plane and its scalability. Monitoring and orchestration of such optical and wireless networks can benefit from the use of software-defined net-

The first two authors contributed equally to this work. Tingjun Chen is with Duke University, USA; Jiakai Yu, Shengxiang Zhu, Steven Santaniello, and Aishik Biswas are with the University of Arizona, USA; Arthur Minakhmetov was with Institut Polytechnique de Paris when this work was carried out, and is now with Alcatel Submarine Networks, France; Craig Gutterman and Gil Zussman are with Columbia University, USA; Michael Sherman is with the University of Chicago, USA; Ivan Seskar is with Rutgers University, USA; Dan Kilper is with Trinity College Dublin, Ireland.

Digital Object Identifier: 10.1109/MNET.006.2100605

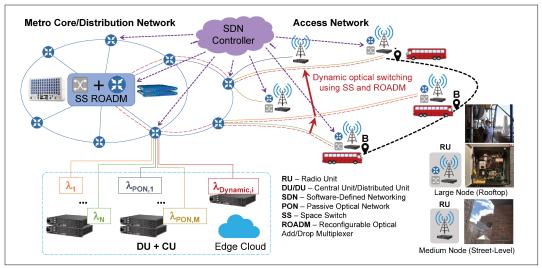


FIGURE 1. Future converged optical-wireless front-/mid-/backhaul (x-haul) networks, where a reconfigurable optical network managed by an SDN controller can support a variety of C-RAN architectures. An example scenario is the WDM/PON C-RAN, where dynamic capacities can be provisioned and allocated to each link depending on the bandwidth and latency requirements.

working (SDN) as a control platform. SDN enables network programmability through logically centralized control, decouples the control plane from the data plane, and abstracts the applications from the hardware [4]. It can largely reduce the complexity of cross-layer, heterogeneous technology systems and provide common map abstraction of network resources in disaggregated systems. The use of machine learning (ML) to improve SDN-based network control also calls for collecting detailed datasets of the key system attributes and functions.

There has been extensive related work on optical networks for wireless and SDN, both theoretical and experimental. While much can be done in small-scale platforms and in lab settings, the optical transmission impairments and their interactions with the control plane require systems of sufficient scale to manifest. Thus, there is a need for experimental city-scale platforms that can enable research on the integration of optical components with new software paradigms (e.g., SDN and network function virtualization [NFV]) and on interworking in real time with heterogeneous wireless and edge computing resources. Such platforms will be beneficial for experimental investigation of novel optical network architectures that can deliver scalable capacity through methods such as amplified rapid optical switching and routing in service to wireless and edge cloud systems, orchestrated by cross-layered and cross-technology SDN-based control planes.

In this article, we present the design and deployment of a field-deployed, remotely accessible, disaggregated and programmable optical-wireless network research platform as part of the city-scale PAWR COSMOS advanced wireless testbed [5], which is being deployed in West Harlem, New York City. Open application programming interfaces (APIs) and programmability across all technology components and protocol layers in COSMOS enable researchers to remotely access the testbed and use it to explore beyond 5G/6G technologies in a real-world environment, including but not limited to edge processing of millimeter-wave baseband signals and analog RoF

communications. We focus on the metropolitan dark fiber-based programmable optical x-haul network as a key infrastructure element of COSMOS, which allows physical layer research throughout the city deployment. Then we present the design and implementation of a dedicated multi-functional Ryu SDN controller for dynamic management of COSMOS' optical network. Important functionalities such as dynamic wavelength channel assignment, network topology reconfiguration, and network monitoring can support a wide range of experiments directly in the optical layer and across higher networking layers.

Finally, we present an example application of the developed SDN controller for monitoring and measurements of the ROADM Erbium-doped fiber amplifier (EDFA) gains, where EDFA gain spectrum data with different channel configurations, collected from deployed equipment in an efficient manner, can be used for ML algorithms. To validate and evaluate the design of COSMOS' optical network and its integration with other testbed components such as software-defined radios (SDRs), we present a pilot experiment focusing on the C-RAN architecture shown in Fig. 1. Particularly, we design and implement a wireless handover (HO) scheme where high traffic volume can be supported by dynamic wavelength allocation via optical switching. In addition, there have been other optical measurements and experiments conducted by testbed users such as that described in [6].

We note that there are few large-scale test-beds for exploration of next-generation network technologies and architectures (e.g., see [3, Section 2, references therein]. For example, Bristol Is Open is a programmable testbed deployed in Bristol, United Kingdom, as a research platform for smart city development, wireless network connectivity, and cloud computing, including optical x-haul elements. The COSMOS network expands on this approach using fully experimental grade physical layer optical and wireless hardware throughout. The ADRENALINE testbed is a circuit-switched laboratory optical testbed designed

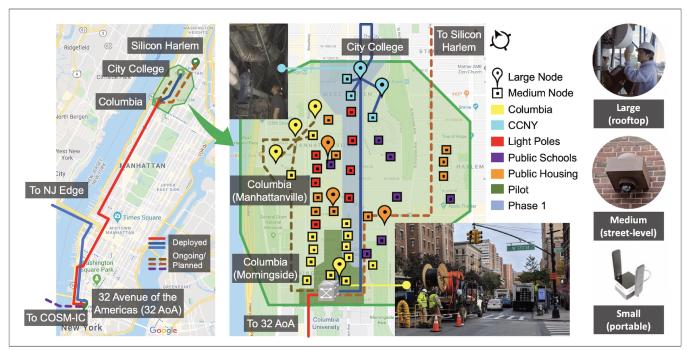


FIGURE 2. Left: COSMOS' envisioned deployment area of ~ 1 mi² in West Harlem, NYC, with deployed and planned dark fiber network, which includes dark fiber between Columbia University and 32 Avenue of the Americas (32 AoA), the City College of New York (CCNY), the Silicon Harlem Gigabit Center, NJ Edge, and COSM-IC; right: the deployment sites of infrastructure radio nodes with different form factors, and the dark fiber interconnections across different sites. The optical switching core of COSMOS, represented by the grey cube, is located in the Computing Research Facilities (CRF) at Columbia University's Computer Science building.

for experimental research on large-scale optical transport networks and distributed edge computing. Generally, these testbeds are limited in the scale and research capability of the optical networks that are connected to the wireless and data infrastructure. Optical testbeds usually surrender research capability as they are integrated with more wireless and data equipment in real-world environments. The availability of dark fiber in NYC and recent advances in disaggregated optical components have enabled COSMOS to provide advanced research capability in the real-world environment with similar research-capable radio and edge cloud computing at scale.

THE PAWR COSMOS TESTBED

COSMOS [5] is a city-scale programmable testbed designed and being deployed to facilitate research and experimentation with advanced wireless and optical technologies in real-world scenarios. Figure 2 shows the full planned deployment, in which a metro dark fiber network and radio nodes (at sub-6 GHz/millimeter-wave frequencies) will be deployed in a dense urban area of $\sim 1 \text{ mi}^2$ in the Harlem neighborhood of Manhattan, in partnership with the NYC CTO, City College of New York (CCNY), Silicon Harlem, and the local community. The envisioned deployment includes 9 large nodes (macrocellular base stations), installed on rooftops, ~40 medium nodes (access points or small cells), installed at street level (building side and lightpoles), and ~200 small (near-portable) nodes, used as fixed or mobile devices throughout the testbed.

A unique feature of COSMOS' optical network is the dark fiber network deployed throughout Manhattan, which allows for various converged optical-wireless experiments in a metropolitan

setting. This dark fiber network, shown in Fig. 2, includes the following.

Deployed:

- A dark fiber ring (~14 m round-trip distance) between the COSMOS optical switching core, located in the computing research facilities (CRF) at Columbia University, and co-location site at 32 Avenue of the Americas (32 AoA)
- Dark fiber between the COSMOS optical switching core and CCNY (~1 mi)
- 100 Gb/s connections between 32 AoA and NJ Edge.

Ongoing:

- Dark fiber between the COSMOS optical switching core and the Silicon Harlem Gigabit Center (~1.8 mi) at 2785 Frederick Douglass Blvd. in Central Harlem
- 100 Gb/s connections between 32 AoA and international collaborative optical and edge cloud testbed infrastructure (e.g., FABRIC in the United States and CONNECT in Ireland) as part of the ongoing NSF COSMOS Interconnecting Continents (COSM-IC) project.

Planned: Dark fiber between the COSMOS optical switching core and the Columbia Manhattanville campus (~1.1 mi).

COSMOS' programmable optical transport network consists of mostly 100 Gb/s+ fiber and freespace optical (FSO) technology, and integrates advanced optical switching technology based on WDM switch fabrics and RoF interfaces to achieve ultra-low-latency connections between different sites and resources. The architecture design of COSMOS' optical network makes use of wavelength-routed switching and space switching in a WDM network to provide two important capabilities: flexible experimentation and network topolo-

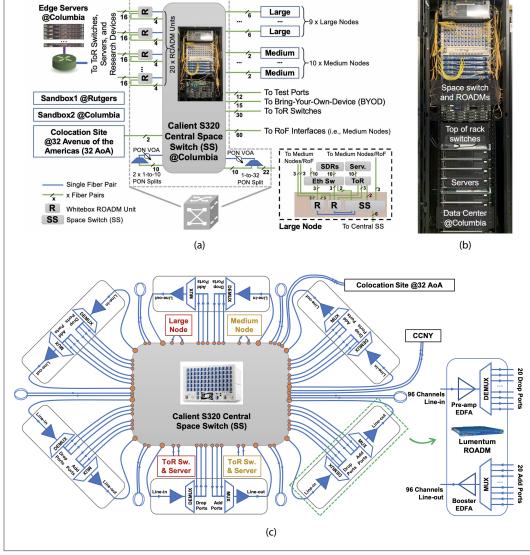


FIGURE 3. a) The envisioned core optical switching architecture of the COSMOS testbed and the switching architecture of a large node in the bottom inset; b)–c) the current deployment and configuration of COSMOS' optical network, where the main components in the optical layer include one Calient S320 central space switch (SS) and six Lumentum ROADM units.

gy reconfiguration of large numbers of radio and computing resource connections, and *multi-lay-ered optical x-haul networking* for experimentation with novel optical devices, systems, architectures, and SDN/NFV optical control planes.

Figure 3a shows the envisioned COSMOS core optical switching architecture, where a Calient $$320 320 \times 320 \text{ space switch (SS) forms the core}$ of COSMOS' optical network in the Columbia CRF. The central S320 SS is connected to smaller 16 \times 16 SSs at each of the large nodes. These SSs allow for remote and automated re-fibering of connections and devices throughout the testbed. WDM is provided by the Lumentum whitebox ROADM units connected to the SSs. Other devices attached to the S320 SS include splitters for PONs, top-ofrack (ToR) switches, test equipment, and experimental hardware. Using these capabilities, the fiber pairs between any two devices (e.g., a radio node and an edge server) can be configured for combinations of point-to-point, PON, and ROADM/ WDM networks. The ROADM units also support flexible grid assignments, where the channel start/stop frequency can be set with minimum granularity of 6.25 GHz. A 96 \times 50 GHz channel C-band configuration is used as the default setting. A range of different interfaces are available for use within the C-band. For DWDM applications, full band tunable long reach DWDM 10 Gb/s on-off keying (OOK) transceivers are available. Short reach fixed and uncooled sources up to 100 Gb/s can be used on dark fibers or with suitably wide channels configured in the ROADMs. Selected radio nodes will also be equipped with analog RoF capabilities in addition to the digital RoF used throughout.

The central \$320 SS will have fiber connections to different sites, whose tentative locations are shown in Fig. 2. Figure 3c shows the *current* optical network connections, the configuration of the testbed, and the main components in the optical layer. In particular, the \$320 SS has connections to six ROADM units (each with two unidirectional EDFAs and WSSs), three ToR switches, one large node, two medium nodes, several edge

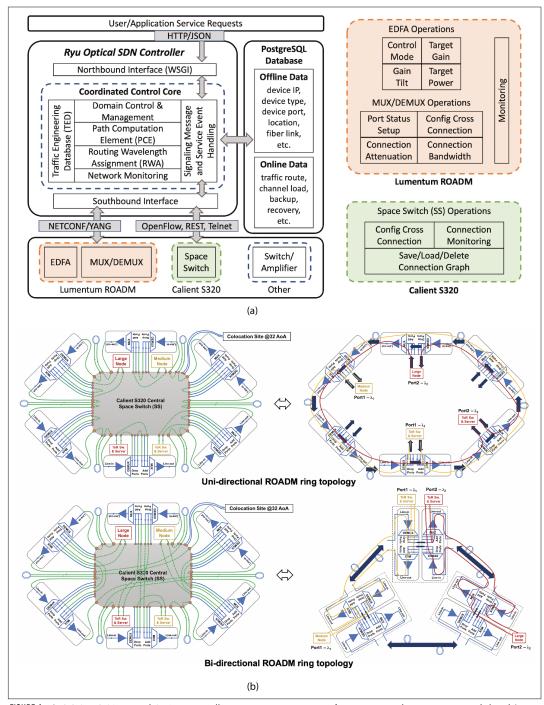


FIGURE 4. a) COSMOS' optical SDN controller: main components, functions, and operation modules; b) two example optical ring topologies supported by the current COSMOS testbed deployment, using the SDN controller, where the red and yellow paths represent the optical signal paths between large/medium nodes and servers on different wavelengths (λ_1 and λ_2).

servers, the co-location site at 32 AoA and CCNY, as well as 12 fiber spools and other devices (e.g., comb source, optical splitters and attenuators, and optical telemetry). A fully transparent x-haul optical transport network can be configured in the current setup using the SS and ROADM units. Additional ROADM units are currently under deployment, and 100/200 Gb/s coherent PM-QPSK/PM-16QAM interfaces will be installed in the near future.

SDN-BASED OPTICAL CONTROL PLANE

COSMOS' optical network is composed of open and whitebox components, allowing for complete freedom in configuring and controlling the network. While COSMOS' optical transport layer will allow advanced users to install their own control and management software, we developed a default SDN controller for managing the testbed's optical resources. Figure 4a shows the control scheme and functional modules of COSMOS' optical SDN controller, which is an enhanced version of the NTT Ryu controller presented in our previous work [7]. Its main components include a

northbound user application web server, a central coordinated control core, a southbound device control agent, and an SQL-based database. In order to support dynamic and flexible network configuration with ROADM units and SSs, the SDN controller is integrated with multiple functional modules, such as a PostgreSQL database, several southbound communication protocols, and space switching control strategy.

Using the SDN controller, the SSs and ROADM units can support reconfigurable optical x-haul connections between wireless devices (e.g., radio nodes) and servers (e.g., edge cloud) for the top-layer user/application requests. In particular, users can send various service requests to the SDN controller including lightpath setup/teardown, optical wavelength channel status monitoring, optical amplifier adjustments, and optical layer topology reconfiguration. These requests are first handled by the northbound web server gateway interface (WSGI), which uses the HTTP protocol and JSON data format for communication between the users and the SDN controller. The service event control and management system coordinates with the traffic engineering database (TED), network monitoring module, routing and wavelength assignment (RWA) algorithm, and path computation element to serve users' requests with available network resources. The southbound agent communicates with the optical physical equipment via various protocols, such as NETCONF, OpenFlow, REST, Telnet, and TL1. Finally, the SDN controller generates an operation result for each request service.

In the current configuration and setup of COS-MOS' optical network (Fig. 3c), two main optical physical layer operations are handled by the SDN controller: flexible traffic lightpath connection operation and dynamic optical network topology reconfiguration. Accordingly, optical network resources can be utilized easily to support various optical x-haul and/or wireless experiments. Figure 4b shows two example optical network topologies.

A 6-Node Unidirectional ROADM Ring Network: Each ROADM unit is connected to another ROADM unit via unidirectional line ports with additional fiber spools by connecting corresponding cross-connection ports of the S320 SS, and a unidirectional pass-through link is established inside each ROADM unit by connecting a drop port to an add port. Two connections between large/medium nodes and servers support data transmission using two different optical wavelengths. Asymmetric round-trip transmission is applied in this unidirectional ring network where the send and receive paths have different numbers of hops. This can be used to maximize the number of hops and propagation distances for transmission experiments.

A 3-Node Bidirectional ROADM Ring Network: A high-degree ROADM node can be formed by establishing two-way pass-through links between multiple ROADM units via add/drop ports. In this example, each 2-degree ROADM node is composed of two ROADM units. Line ports are used to build duplex connections with additional fiber spools between these colorless 2-degree ROADM nodes. This full bidirectional ROADM ring is formed by two opposite unidirectional rings, where the connections between large/medium nodes and servers have symmetric round-trip lightpaths for signal transmission. Using high-degree ROADM

Using the SDN controller, the SSs and ROADM units can support reconfigurable optical x-haul connections between wireless devices (e.g., radio nodes) and servers (e.g., edge cloud) for the top-layer user/application requests.

nodes, more complex network topologies (e.g., star/mesh networks) can be configured in COS-MOS' optical layer.

Using the deployed Lumentum ROADM units, dark fiber networks connecting different sites including 32 AoA and CCNY with ~14 mi and ~1 mi distances, respectively (Fig. 2), and a spooled fiber plant with 12 fiber spools (~100 mi total distance), a large number of city-scale metropolitan optical network topologies can be customized by users for novel experiments. In addition, a wider range of topologies can be realized by routing the 96 available wavelengths transparently between various nodes in combination with the dark fiber network and spooled fiber plant. More ROADM units and fiber spools will be deployed in COS-MOS in the near future. These networks can then feed into or be extended to the RAN, creating a rich experimental environment for x-haul networks at city-scale.

NETWORK MONITORING AND DATA COLLECTION: AMPLIFIER CHARACTERIZATION

In transparent optical networks, optical signals traverse through network nodes without optical-electrical-optical conversion, and therefore, it is important to evaluate the quality of transmission (QoT) before an optical channel is provisioned for service. Since optical channels share the same fiber strand, the add/drop operations of optical wavelength channels can affect other channels through effects such as cross-talk, fiber nonlinearity, and channel power excursions. Due to the complexity of these systems, the use of ML algorithms for QoT estimation and functions (e.g., fault management) has recently gained extensive attention [8, 9]. COSMOS' optical network provides a unique opportunity to collect detailed data for the development of such algorithms. In this section, we present an example application of COSMOS' optical SDN controller for characterizing the ROADM EDFAs, whose gain spectrum usually depends on the channel configuration and input power. We develop methods and SDN control strategy for automatic EDFA gain spectrum data collection from the optical layer. Such data can be used to develop new optical channel provisioning and planning tools including QoT estimation and impairment-aware RWA.

In COSMOS' optical network, each ROADM unit has an internal optical channel monitor (OCM) that captures wavelength spectrum information. The Ryu SDN controller communicates with the OCM through the southbound interface and retrieves all ROADM data via NETCONF/YANG (Fig. 4a), including the EDFA monitoring data, total input/output power, VOA setting, loss, gain mode, gain value, and gain tilt. In particular, when the SDN controller receives a monitoring request for an ROADM unit, the retrieved data from the OCM is converted into JSON data format and sent back to users through the north-

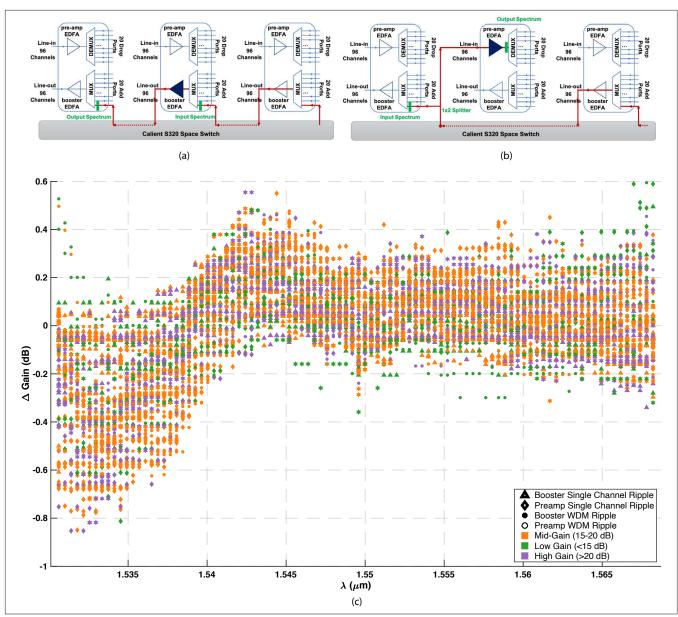


FIGURE 5. Setups for booster and pre-amp EDFA gain spectrum measurements in a Lumentum ROADM unit: a) the booster EDFA characterization; b) the pre-amp EDFA characterization; c) measured ROADM booster and pre-amp EDFA gain ripple spectrum (ΔGain) with single/WDM channels and different gain settings.

bound WSGI API. Each ROADM unit has two EDFA modules (Fig. 3c): a booster EDFA after the MUX WSS, and a pre-amp EDFA before the DEMUX WSS. Since the OCM does not have specific gain spectrum or power spectrum monitoring for the EDFA only, the wavelength-dependent gain measurements of the two EDFA modules require different data acquisition procedures. Figure 5a shows the booster EDFA gain spectrum measurement setup, where the booster EDFA in the middle ROADM is the target EDFA. The channel input power spectrum is measured at the MUX WSS module in the same ROADM, and the channel output power spectrum is captured at the next hop MUX WSS module. The transmission loss in the SS can be retrieved by the SDN controller. Thus, the target booster EDFA gain spectrum can be calculated from the corresponding measurements. Similarly, Fig. 5b shows the preamp EDFA gain spectrum measurement setup,

where the channel output power spectrum of the target pre-amp EDFA in the middle ROADM can be directly measured using its internal DEMUX WSS. However, the channel input power spectrum cannot be obtained either from its own ROADM WSS OCM, since there is no OCM module before the pre-amp EDFA, or from the previous hop ROADM because the booster EDFA output power is directly fed into the target pre-amp EDFA. Thus, the signals from the previous hop are directly fed into an available MUX port of the next-hop ROADM unit through an optical splitter, and the input power spectrum can be retrieved using the MUX OCM of the next ROADM unit with appropriate calibration.

We consider two measures for wavelength-dependent gain: single channel ripple and WDM ripple, where the single channel ripple is measured as each channel is individually activated, and the WDM ripple is recorded in a single measurement

with all 96 channels on. We developed a customized script, including repeated SDN controller requests of traffic lightpath operations, space switching operations, and received data recordings, which measures the wavelength-dependent gain of both EDFA modules in COSMOS' optical transport layer. Using this script, more than 8000 datasets corresponding to different operating conditions were generated for each EDFA in the 6 ROADM units deployed in COSMOS (Fig. 3c). Figure 5c presents both the single and WDM channel gain ripple measurement results of all 12 EDFAs operating at different gain levels. It can be seen that the gain ripple depends not only on the specific EDFA unit, but also on the gain setting and channel loading configuration. These measurements provide valuable datasets for the development of ML algorithms [10] to improve both component- and system-level performance in optical networks, as well as for building realistic hardware models for a digital twin of large-scale optical networks [11]. Additional data will be recorded as new ROADM units are deployed and as new topologies are realized over the life of the testbed, and all collected datasets will be made publicly available to the community. Users can also conduct a broader range of measurements and experiments using COSMOS' optical network and SDN controller depending on their research requirements.

PILOT EXPERIMENT: C-RAN WIRELESS HANDOVER USING OPTICAL SWITCHING

In this section, we consider the application of optical-switching-based wireless handover (HO) between RUs connected to a mobile UE that requires exceptionally large capacity with deterministic low latency, and propose a scheme of dynamic capacity allocation of optical/wireless links to support varying network bandwidth requirements. An example application scenario is shown in Figs. 1 and 6a, where vehicles move in a dense urban area and passengers use high data rate applications (e.g., real-time multi-view vehicle video analytics). When the UE is moving from the coverage area of one RU to another, the capacity of the underlying optical fronthaul link needs to be dynamically and seamlessly re-allocated. To avoid provisioning high capacity throughout the network, a channel on a dedicated wavelength can be optically switched from one link to the other link using either the optical SSs or the ROADM units. Using optical networks to facilitate wireless HO has been investigated before (e.g., in [12]), and the proposed approaches are candidates for experimental evaluation in COSMOS.

In [13], we implemented and experimentally evaluated optical-switching-based wireless HO in C-RANs using COSMOS' optical and radio resources together with the SDN controller, which is well suited for the study of such multi-technology SDN control scenarios. In this pilot experiment, SDN control was implemented for the edge cloud including the DU/CU along with the optical switching, and the HO decision on optical switching from the source RU to the target RU was made based on the wireless signal power threshold. Figure 6b shows the signaling flows of two considered HO modes:

In this pilot experiment, SDN control was implemented for the edge cloud including the DU/CU along with the optical switching, and the HO decision on optical switching from the source RU to the target RU was made based on the wireless signal power threshold.

- Hard HO with non-zero downtime ("break before make"): The SDN controller can assign the UE a single wavelength that "follows" the UE from the source RU to the target RU. The switching consists of changing the SS configuration, always keeping one port connected to the link leading to the edge cloud and changing ports to the link leading to the RUs.
- Soft HO with zero downtime ("make before break"): This uses at least two wavelengths or independent fiber channels at the time of HO, and the new connections from the edge cloud to an RU are made before the old connections are broken using additional network resources.

Figure 6c shows the experimental setup involving two RUs located on two sides of the Columbia CEPSR Building, each emulated using two USRP N210 SDRs with one assigned a fixed capacity and one assigned a dynamic capacity. The link with fixed capacity between the RU and the edge cloud is always live, whereas the link with dynamic capacity between the RU and the edge cloud can be dynamically allocated via optical switching. For simplicity, we consider only upstream traffic in this experiment, which does not limit the validity of our observations. All USRPs are connected to a Pica8 ToR switch via 1 Gb/s Ethernet, which has four bidirectional tunable C-Band 10 Gb/s transceivers connected to the S320 SS located in the Columbia CRF. Specifically, three ports are used for hard HO, and the fourth port is required for soft HO. The S320 SS routes a lightpath from the Pica8 ToR switch toward another ToR switch that has a 25 Gb/s connection to edge cloud servers. To avoid switching transients, we use direct point-to-point WDM connections over the SS as opposed to the ROADM units. Moreover, the SS is used for switching due to its faster SDN controlled switching time. We implemented the edge cloud functionalities including the SDN controller/ agent and remote I/Q signal processing received by the RUs using the GNU Radio software on a COSMOS edge server. We used a USRP B205-mini SDR as a UE that transmits a continuous wave tone at 900 MHz carrier frequency. The RUs stream the received I/Q samples to the edge cloud server, and the received signal strength is measured at 0.1 ms intervals. A switching and HO event is triggered if the received signal strength becomes lower than a threshold (-57 dBm), and the dynamic capacity is switched from the source RU to the target RU (more details can be found in [13]).

We performed measurements of the *HO time* (i.e., the duration between the triggering and completion of the HO event) and *downtime* (i.e., duration in which no RU is connected to the edge cloud). A total number of 29 trials were conducted for the hard HO and soft HO schemes using 10 Gb/s optical links between the edge cloud and USRPs. We also conducted experiments with 1 Gb/s optical links, where the HO conditions were emulated in software. Our experimental

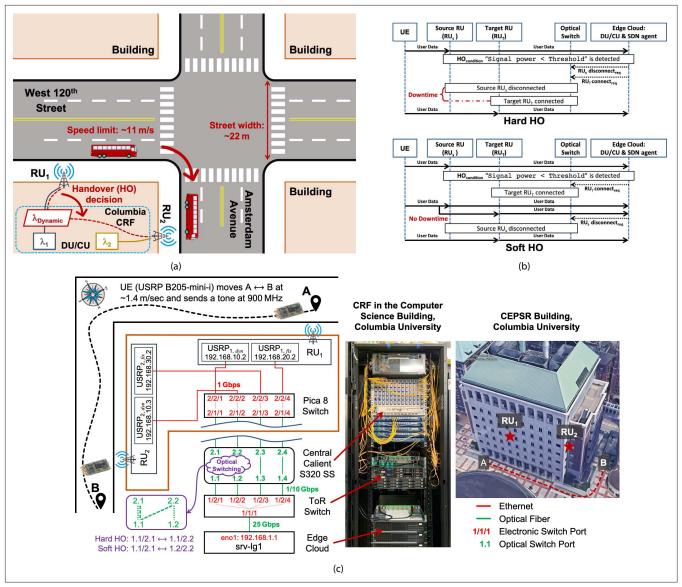


FIGURE 6. a) Illustration of the C-RAN handover (HO) scheme at the COSMOS Pilot deployment site, where a vehicle taking a turn at an intersection is served by two radio units (RUs) through dynamic optical switching and wavelength assignment; b) the signaling flows of the two HO modes: hard HO ("break before make") and soft HO ("make before break"); c) setup of the C-RAN HO experiment in the COSMOS testbed, where the user equipment moves between two RUs and invokes HO based on the HO signaling flows shown in a).

results showed that the average switching time for hard HO is 0.96/1.83 s (standard deviation: 0.15/0.24 s) with 1/10 Gb/s links, respectively, where the corresponding downtime is 0.91/1.65 s (standard deviation: 0.15/0.32 s). Similarly, the average switching time for soft HO is 0.72/1.61 s (standard deviation: 0.08/0.21 s) with 1/10 Gb/s links, respectively, with zero downtime. The HO switching time is limited by the time it takes for the optical transceivers to establish the link connections, which is on the order of seconds [14]. These results were confirmed by WireShark-based measurements of the optical link recovery time. The time difference between 1 Gb/s and 10 Gb/s is due to the fact that the 1 Gb/s links are fixed rate with no rate negotiation phase during link establishment. We note that the long HO times can be reduced by using burst-mode transceivers, and we will also consider optical switching using the ROADM units in future work.

CONCLUSIONS

In this article, we present the design and deployment of an open access optical-wireless network research platform as the key technological component of the PAWR COSMOS testbed. COS-MOS' programmable optical infrastructure consists of open and whitebox optical space switches, ROADM units, and a dark fiber network, orchestrated by an SDN control plane, allowing for the configuration of a variety of network topologies in the optical physical layer. We also present the use case of the SDN controller for network monitoring and automated data collected, as well as a pilot experiment on wireless handover supported by optical switching. Together with the radio and compute resources available in the testbed, COS-MOS' optical network can support a broad range of C-RAN and convergent optical-wireless experiments in a real metropolitan setting.

ACKNOWLEDGMENTS

Partial and preliminary results were presented in [5,13,15]. This research was supported in part by NSF grants CNS-1737453 and CNS-1827923, NSF-BSF grant CNS-1910757, OAC-2029295, CNS-2128638, CNS-2112562, Science Foundation Ireland under Grant #13/RC/2077_P2, the IMT/Columbia University's COMITI exchange program, an AT&T VURI Award, a Google Research Scholar Award, and an IBM Academic Award.

REFERENCES

- [1] T. Pfeiffer, "Next Generation Mobile Fronthaul and Midhaul Architectures," *IEEE/OSA J. Optical Commun. Networking*, vol. 7, no. 11, 2015, pp. B38–45.
- [2] M. Ruffini, "Multidimensional Convergence in Future 5G Networks," IEEE/OSA J. Lightwave Technol., vol. 35, no. 3, 2016, pp. 535–49.
- [3] D. M. Marom et al., "Survey of Photonic Switching Architectures and Technologies in Support of Spatially and Spectrally Flexible Optical Networking," *IEEE/OSA J. Optical Commun. Networking*, vol. 9, no. 1, 2017, pp. 1–26.
- [4] H. Kim and N. Feamster, "Improving Network Management with Software Defined Networking," *IEEE Commun. Mag.*, vol. 51, no. 2, Feb. 2013, pp. 114–19.
- vol. 51, no. 2, Feb. 2013, pp. 114–19.
 [5] D. Raychaudhuri et al., "Challenge: COSMOS: A City-Scale Programmable Testbed for Experimentation with Advanced Wireless," Proc. ACM MobiCom '20, 2020.
- [6] K. Guan et al., "Efficient Classification of Polarization Events Based on Field Measurements," Proc. IEEE/OSA OFC '20, 2020, Th3D.7.
- [7] Y. Li et al., "tSDX: Enabling Impairment-Aware Alloptical Inter-Domain Exchange," *IEEE/OSA J. Lightwave Technol.*, vol. 36, no. 1, 2018, pp. 142–54.
- vol. 36, no. 1, 2018, pp. 142–54.
 [8] A. D'Amico et al., "Using Machine Learning in an Open Optical Line System Controller," IEEE/OSA J. Optical Commun. Networking, vol. 12, no. 6, 2020, pp. C1–11.
 [9] A. Mahajan et al., "Modeling EDFA Gain Ripple and Filter Pen-
- [9] A. Mahajan et al., "Modeling EDFA Gain Ripple and Filter Pen alties with Machine Learning for Accurate Qot Estimation," J Lightwave Technol., vol. 38, no. 9, 2020, pp. 2616–29.
- [10] S. Zhu et al., "Hybrid Machine Learning EDFA Model," Proc. IEEE/OSA OFC '20, 2020, T4B.4.
- [11] D. Wang et al., "The Role of Digital Twin in Optical Communication: Fault Management, Hardware Configuration, and Transmission Simulation," *IEEE Commun. Mag.*, vol. 59, no. 1, Jan. 2021, pp. 133–39.
 [12] X. Wang et al., "Handover Reduction via Mobility Predictions."
- [12] X. Wang et al., "Handover Reduction via Mobility Prediction-Based VPON Formation in Optical-Access-Enabled Cloud-RAN," Proc. IEEE/OSA OFC '15, 2015.
- [13] A. Minakhmetov et al., "Experiments on Cloud-RAN Wireless Handover Using Optical Switching in a Dense Urban Testbed," Proc. IEEE/OSA OFC '20, 2020, Th2A.25.
- [14] P. Cibula, "Mechanisms Influencing BASE-T 1Gb/10Gb Time-to-Link," IEEE P802.3bz 2.5/5GBASE-T, 2015.
- [15] J. Yu et al., "Dual Use SDN Controller for Management and Experimentation in a Field Deployed Testbed," Proc. IEEE/ OSA OFC '20, 2020, T3J.3.

BIOGRAPHIES

TINGJUN CHEN received his Ph.D. degree in electrical engineering from Columbia University in 2020. Between 2020 and 2021 he was a postdoctoral associate at Yale University. Since Fall 2021 he has been with Duke University, where he is currently an assistant professor of electrical and computer engineering. His research interests are in the areas of networking and communications with a focus on next-generation wireless and optical networks. He received the IBM Academic Award, Google Research Scholar Award, Columbia Engineering Morton B. Friedman Memorial Prize for Excellence, and Facebook Fellowship. He also received the ACM CoNEXT '16 Best Paper Award and the ACM SIGMOBILE Dissertation Award runner-up.

JIAKAI YU received his Ph.D. degree in electrical and computer engineering from the University of Arizona in 2021, his M.S. degree in electrical engineering from the State University of New York at Buffalo in 2015, and a degree at Physics College from Nankai University, Tianjin, China, in 2013. His research interests include software-defined networking, data center/campus/transport networks, and Al-based intelligent networks.

ARTHUR MINAKHMETOV graduated with an M.S. degree in electronics, electrical energy, and automatics from Paris-Saclay University and with an engineering degree from Institut d'Optique Graduate

School in 2016. He obtained his Ph.D. degree in networks, information, and communications from Telecom Paris in 2019. During 2019, he was a visiting scholar at Columbia University, where he contributed to the design and deployment of the COSMOS optical testbed. In 2020, he joined Alcatel Submarine Networks. His research interests include fast switching optical networks, congestion control algorithms, wireless-optical network convergence, optical transmission related signal impairments, and application of machine learning techniques for optical networking problems.

CRAIG GUTTERMAN received his B.Sc. degree in electrical engineering from Rutgers University, New Jersey, in 2012 and his Ph.D in electrical engineering from Columbia University, New York, in 2021. He received the NSF GRFP and the From Data to Solutions NSF IGERT fellowships. His research focuses on improving the performance of future networks and systems by developing machine-learning-based network systems and data-driven network algorithms.

MICHAEL SHERMAN is the infrastructure lead for the Chameleon Testbed, where he works on enabling reproducible computer science research. Prior to this he helped develop the ORBIT and COSMOS testbeds, working on reproducible experiments in wireless and edge computing. His research interests are in the reliability of large systems at all levels: computing, networking, and human interaction. He is a senior scientist in the Department of Computer Science at the University of Chicago and holds a joint appointment at Argonne National Laboratory.

SHENGXIANG ZHU received his Bachelor's degree in optoelectronics from Shenzhen University, China. He received his Ph.D. degree from the Department of Electrical and Computer Engineering, University of Arizona, Tucson. His research interests include coherent optical transmission, optical network control, software-defined networks, and photonics.

STEVEN SANTANIELLO received his M.S. degree in optical sciences from the University of Arizona, and joined the Air Force as an engineer in August 2021. He performed his graduate studies working with Dr. Dan Kilper on data-driven engineering of optical networks, particularly on statistical characterization of EDFAs. He also spent time working on ovarian cancer research with Dr. Jennifer Barton, looking at statistical characterization of fluorescent ovarian tissue images.

AISHIK BISWAS received his M.S. degree in optical sciences from the University of Arizona in 2020, where he was advised by Dr. Daniel C. Kilper. He completed his B.Tech. degree in electronics and instrumentation engineering from the West Bengal University of Technology in 2017. He is currently pursuing a Ph.D. degree in electrical engineering at Montana State University. His research interests include optical communication systems and photonic devices.

IVAN SESKAR is the chief technologist at WINLAB, Rutgers University, responsible for experimental systems and prototyping projects. He is currently the program director for the COSMOS project responsible for the New York City NSF PAWR deployment and was the co-PI and project manager for all three phases of the NSF-supported ORBIT mid-scale testbed project at WINLAB, successfully leading technology development and operations since 2005 and for which the team received the 2008 NSF Alexander Schwarzkopf Prize for Technological Innovation. He is a Co-Chair of the IEEE Future Networks Testbed Working Group and the co-founder and CTO of Upside Wireless Inc.

GIL ZUSSMAN [F] received his Ph.D. degree in electrical engineering from the Technion in 2004 and was a postdoctoral associate at MIT in 2004–2007. Since 2007 he has been with Columbia University where he is currently a professor of electrical engineering and computer science (affiliated). His research interests are in the areas of wireless, mobile, and resilient networks. He is a recipient of the Fulbright Fellowship, DTRA Young Investigator Award, NSF CAREER Award, and two Marie Curie Fellowships. He is a co-recipient of seven best paper awards and is the Columbia PI of the NSF PAWR COSMOS testbed.

DAN KILPER is a professor of future communication networks and SFI CONNECT Centre director at Trinity College Dublin, Ireland. He holds an adjunct faculty appointment at the Columbia University Data Science Institute and the College of Optical Sciences, University of Arizona. He received M.S. (1992) and Ph.D. (1996) degrees in physics from the University of Michigan. From 2000 to 2013, he was a member of technical staff at Bell Labs, and he co-founded Palo Verde Networks. His research is aimed at solving fundamental and real-world problems in communication networks, addressing interdisciplinary challenges for smart cities, sustainability, and digital equity.