

# Design and Experimental Evaluation of Wireless Interface Agents for SDN mmWave Multi-Hop Networks

Olivia R. Thompson, Ethan W. Lee, Ava M. Patel

Olivia R. Thompson and Ethan W. Lee, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA; Ava M. Patel, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

## Abstract

Millimeter wave (mmwave) communications will likely be an enabler for 5G due to its multi-gigabit per second throughput capabilities. Furthermore, mmWave communications will have to be integrated in a new redesigned network required by 5G to fulfill its ambitious targets. In this paper, we present the design and implementation of a management agent for wireless devices deployed in a heterogeneous SDN wireless multi-hop research platform featuring mmwave communications for crosshauling (backhaul and fronthaul) purposes. The performance of the deployed mmwave network, based on the IEEE 802.11ad standard, is measured employing this agent. We measure the downtime in the presence of link up/down events, with obtained response times in the order of 10s-to-100s of milliseconds depending on the case. Furthermore, the TCP performance over the multi-hop 802.11ad mmwave network is also experimentally evaluated. In fact, TCP throughput up to around 800Mbps are obtained for single and multi-hop scenarios despite neighboring links using the same channel. Finally, one can also observe the impact of MTU size on TCP throughput, which may hinder the full exploitation of the mmWave link capacity when combined with other transport technologies, since the advantages of big MTUs (much bigger than the typical 1500 bytes) offered

**Indexed keywords:** Computer Engineering, Advanced Computing, Technology, Open Access

Article History: Received: 16 July 2025 | Accepted: 09 September 2025 | Published: 27 September 2025

mmwave, SDN, 5G, IEEE 802.11ad, integrated management, experimental study, mesh network

## 1 INTRODUCTION

Upcoming 5G networks are addressing ambitious Key Performance Indicators (KPIs) in terms of not only overall/per user capacity and latency, but also, in terms of network control and management. Supporting these requirements allows 5G networks to expand their service scope to vertical industries, such as eHealth, automotive, media, or industry 4.0, hence enriching the telecom ecosystem.

On one hand, 5G redefines the air interface and communications at millimeter wave (mmwave) bands (28-100 GHz frequency range) will likely to be an enabler to achieve the aforementioned objectives due to its multi-gigabit per second capacity thanks to the large chunks of contiguous available bandwidth. On the other hand, 5G networks will also require a redesign of the overall network architecture, where complexity and heterogeneity will be the rule, hence requiring increasing automation and programmability. Furthermore, the borders between traditional network segments (access, transport and core) will likely to be blurred and distributed computing resources will be deployed across the mobile network infrastructure. The consequent research problems of these network scenarios are tackled by EU H2020 5G-PPP phase 2 5GTRANSFORMER project [2]. In this sense, 5G embraces software defined networking (SDN) and network function virtualization (NFV) principles as key-enabling techniques to realize this vision.

In this paper, we focus on applying SDN principles to a transport network research platform integrating backhaul and fronthaul (crosshauling) traffic, which features heterogeneous wireless links, namely mmwave and Wi-Fi [10]. The application of such principles allows a unified and dynamic control and management of integrated 5G multi-technology wireless transport networks [4]. This contrasts with cumbersome, error-prone, and inefficient traditional network management



approaches based on the configuration of proprietary devices [15], which is often done manually.

In particular, our contributions in this paper are as follows. First, we detail the design and implementation of the wireless interface agent (WIA) integrated in our SDN research platform [10]. Its purpose is to manage the wireless transport interfaces of its nodes. The WIA can manage heterogeneous transmission technologies, e.g., mmwave and Wi-Fi, and it is employed to characterize the performance of the deployed mmwave network. Second, we measure the downtime of the mmwave network in the presence of link up/down events triggered with the WIA, with obtained response times in the order of 10s-to-100s of milliseconds depending on the case. Third, the TCP performance over the multi-hop mmwave network is also experimentally evaluated. In fact, TCP throughput up to around 800Mbps are obtained for single and multi-hop scenarios despite neighboring links using the same channel. Finally, one can also observe the impact of Maximum Transmission Unit (MTU) size on TCP throughput, which may hinder the full exploitation of the mmwave link capacity when combined with other transport technologies, since the advantages of big MTUs (much bigger than the typical 1500 bytes) offered by mmwave devices may not be reaped.

The remainder of this paper is organized as follows. We summarize related work in Section 2 with respect to management of wireless SDN transport networks and characterization of experimental mmwave networks. Section 3 presents the developed agent to manage mmwave devices in an SDN-enabled wireless mesh network. Section 4 and Section 5 present the scenario under assessment and its characterization, respectively, which has been obtained using the presented management agent. Finally, we conclude our work in Section 6.

## 2 BACKGROUND AND RELATED WORK

### 2.1 Management agents for wireless SDN networks

Currently, the control and the management of wireless transport equipment is mostly based on proprietary solutions using different versions of Simple Network Management Protocol (SNMP) protocol or cumbersome command line interface (CLI), limiting automation and requiring specialized network engineers [15].

In order to overcome these limitations, the Open Networking Foundation (ONF) created the Wireless Transport SDN group to apply SDN principles to wireless transport networks. In their first Proof of Concept (PoC), back in 2015, they considered extensions of the Open Flow Protocol (OF) to handle not only the control but also the management of wireless transport equipment. This approach has also been followed by other research work, such as [3]. However, the Wireless Transport SDN group of the ONF realised that common information models of network devices are required in order to achieve a more automated and unified management of wireless devices, decoupling the management from the control operations (forwarding behavior).

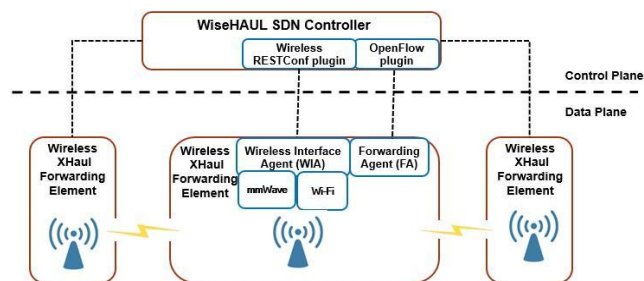
Then, the following PoCs of the Wireless Transport SDN group of ONF (the last one carried out in August 2017 [8]), demonstrated the capabilities and benefits from utilizing a common Information Model (ONF TR-532) for multi-vendor management of microwave transport network devices. This information model is exchanged by means of the Network Configuration (NETCONF) protocol, a general-purpose management protocol, which uses YANG modeling language to implement the data model. At the network device, an adapter, what ONF calls mediator, is used for translating the NETCONF/YANG information model to/from the existing proprietary management procedures of each vendor's device. This approach is similar to the one we present in Section 3. However, our approach aims to manage devices using different technologies, namely mmwave and Wi-Fi, and we use RESTConf interface to exchange the information, allowing a simplified transaction model with respect to NETCONF, hence allowing for faster prototyping.

### 2.2 Characterization of experimental

#### mmwave networks in the 60GHz band

The interest of mmwave technology to enable 5G network system design is evident due to the availability of large contiguous spectrum availability to deliver high speed data rates. Between the mmwave bands, the V-Band (60GHz frequency band) is of special interest because it is a worldwide licence free spectrum. This band is used in the IEEE802.11ad/WiGig standard [9], a promising solution for mmwave radio communications, which could be used for mobile crosshauling purposes in 5G ultra dense networks. However, there is still a shortage of experimental

evaluation of IEEE 802.11ad networks. Some examples on the characterization of real singlehop IEEE 802.11ad mmwave networks are [6], [11] and [12], which consider different scenarios and conditions such as outdoor picocells, indoor office environment or the effect of co-channel and adjacent channel interference. However, these works focus on single-hop link-level behavior characterization and do not consider interface management of mmwave nodes in a SDN scenario. Closer to our work, we find [5], where resiliency mechanisms based on the OF protocol in a multi-hop SDN-enabled IEEE 802.11ad mmwave mesh testbed are assessed with UDP flows. Contrarily, this paper focuses more on the management part of mmwave devices in an SDN-enabled network and how the developed agent is employed to characterize the performance of the



Paper Session X

Figure 1: WiseHAUL SDN platform architecture

deployed mmwave involving control and data plane SDN architectural layers. It is worth highlighting that for the latter, TCP flows are used, as opposed to UDP in previous work [5].

### 3 WIRELESS INTERFACE AGENT FOR SDN WIRELESS DEVICES

Figure 1 presents the architecture of WiseHAUL [10], an SDN-empowered platform combining heterogeneous wireless technologies (mmwave and Wi-Fi) to conduct research and experimentation on SDN/NFV for 5G mobile transport networks. This section deals with the Wireless CrossHaul Forwarding Elements (WXFE) present in the WiseHAUL platform, focusing on the developed Wireless Interface Agent (WIA). As we can see in Figure 1, the SDN agent embedded in the WXFE has two components. This separation obeys to the logic of decoupling the control

operations (forwarding behavior) from the management operations (configuration and administration of network elements). On one hand, we have the Forwarding Agent (FA), which is the OF protocol [7] agent provided by the instance of the open source software switch running in the transport node. This agent oversees forwarding actions and it is technology-agnostic of the underlying wireless transport technology (e.g., mmwave, Wi-Fi) since the matching headers taken into account to establish flow rules present an equivalent frame format. On the other hand, each wireless transport node counts with its own WIA, which exposes an Application Programming Interface (API) to the WiseHAUL SDN controller and upper application layers allowing them to configure the underlying wireless network devices.

The WIA is able to integrate interfaces of heterogeneous wireless technologies, by means of an appropriate information model defined using the YANG modeling language. Then, the WIA is in charge of mapping the request to the specific configuration command of each interface of the transport node. The WiseHAUL SDN controller and the WXFE communicate via a RESTConf interface over HTTP requests, where a uniform resource identifier (URI) identifies every exchanged resource. The client runs in

mmNets '18, October 29, 2018, New Delhi, India the WiseHAUL SDN controller and the server runs in each WXFE.

This WIA approach (YANG-RESTConf) provides flexibility and scalability capabilities to the solution if we compare to custom extensions of OF protocol, since the addition of new features is handled easily by extending the information model and defining its associated URIs. An extension of OF protocol, initially proposed in the 1st PoC of the Wireless Transport Group of ONF and in [3], requires the definition of new structures and parameters within the OF protocol structures/messages and the corresponding inclusion of code into the internals of the used open source software switch and the SDN controller to appropriately parse new defined OF messages, hence limiting the scope and scalability of the solution. Furthermore, the use of RESTConf over NETCONF simplifies the interface, easing the prototyping, since the set of NETCONF operations can be implemented with the reduced set of HTTP methods.

In addition to configuration parameters, the WIA considers asynchronous notifications of network events

following a publish/subscribe pattern, which could be used to trigger actions at the SDN controller level based on monitored events at the network nodes.

Table 1: Current parameters defined at the YANG model for wireless interfaces

Parameter	Config	Monitor	Inventory
Supported Channels			X
Channel central frequency	X	X	
Supported Chan Bandwidth			X
Channel Bandwidth	X	X	
Supported MCS			X
Current MCS	X	X	
Transmission Power	X	X	
Max Tx power per channel			X
Link Status	X	X	

Table 1 summarizes the set of currently supported parameters by the WIA, based on capabilities of available hardware based on Wi-Fi (Compex WLE900VX IEEE 802.11ac) and mmwave (TensorCom TC60G-USB3.0 EVK IEEE 802.11ad) technologies. The available mmwave devices count with reduced configuration capabilities with respect to the considered Wi-Fi device because it is a pre-commercial device. In particular, current capabilities of the mmwave device drivers are reduced to read/configure its modulation and coding scheme (MCS) and link status, and read its central frequency.

#### 4 EXPERIMENTAL SETUP

We have configured a four mmwave node transport network topology using five nodes of the WiseHAUL platform, where four nodes are WXFES and the remaining one is running the WiseHAUL SDN controller. The used setup is depicted in

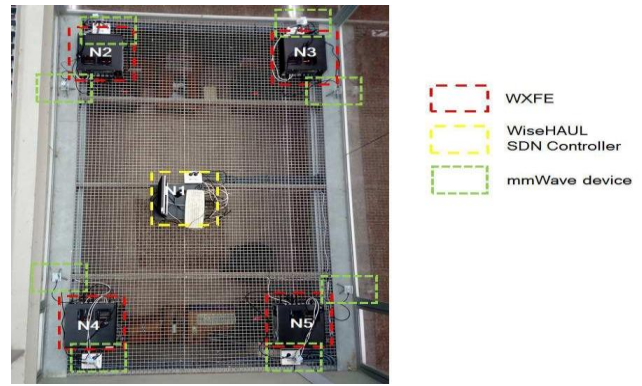


Figure 2: Experimental mmWave network setup under assessment

Figure 2. The distance between N4-N2 and N5-N3 WXFES is four metres, while the distance between N4-N5 and N2-N3 WXFES is two metres.

Each node is based on an Intel Core i7 processor (12-Cores at 3.3GHz x86 CPU) with 32GB and 1TB of hard disk, running the Ubuntu 16.04 LTS distribution. Each node is equipped with three Compex WLE900VX IEEE 802.11ac Wi-Fi cards running the *ath10k* driver and the firmware provided by Candela Technologies [1], which is necessary to establish IBSS connections with adjacent nodes. They also count with two Gigabit Ethernet ports, one dedicated to testbed management purposes and the other one to set an out-of-band wired SDN control plane channel. WXFES has also USB 3.0 slots where up to two TensorCom TC60G-USB3.0 EVK devices are plugged to provide the platform with IEEE 802.11ad mmwave capabilities. The mmwave devices are mounted as network interfaces when loading its accompanying drivers in the kernel. In order to establish the mmwave link, the TensorCom devices are configured to operate as PBSS Control Point (PCP) mode on one side and as station on the other side. This mode configuration is done by means of control scripts also provided by TensorCom Inc. WXFES use OpenvSwitch (OVS) with version 2.7.0 as software switch. In these switches, the different wireless interfaces, both mmwave and Wi-Fi, are added as OVS ports for transport purposes. An additional port using linux namespaces is added to the OVS switch to generate/receive traffic.

#### 5 EXPERIMENTAL RESULTS

In this section, we present the performance characterization of the described experimental setup.

This is done through two case studies. The first case study provides results showing the experienced downtime after a change in the status of the mmwave link, while the second one studies the experienced TCP throughput in single and multi-hop mmwave setups when varying MCS and MTU. The experiments are executed with the use of the described WIA, triggered from the Graphical User Interface (GUI) of the WiseHAUL SDN controller. The reported values represented in the following figures show the statistical distribution of the corresponding experiment from the 25<sup>th</sup> to the 75<sup>th</sup> percentiles, and the whiskers from the 5<sup>th</sup> to the 95<sup>th</sup> percentiles. The marker represents the average value.

### 5.1 Case Study I: Setting the link status

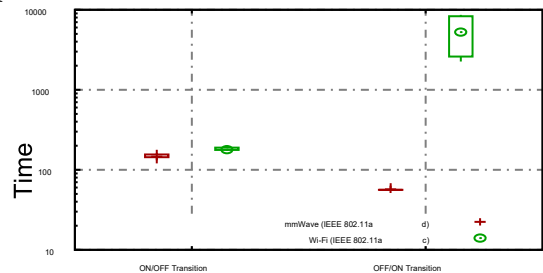
The following results show the experienced downtime in a mmwave link when there is a change in its status. This change can happen unintentionally, due to a failure in the device, or intentionally through a management decision.

This study is possible thanks to the recovery possibilities provided by the WiseHAUL SDN controller [4] to react in front of network events, such as a change in the network topology. After a link event (either down or up) is reported, the WiseHAUL SDN controller checks the affected paths, computes, deletes and reinstalls alternative paths (if possible). In this case, a "break-before-make" strategy is implemented. Next, we compare the downtime of a mmwave link with the one experienced by the Wi-Fi IEEE 802.11ac link. In this way, we show how the WiseHAUL SDN controller interacts with the WIA and the FA of the different WXFЕ depicted in Figure 1 and the capabilities of the WIA to simultaneously manage heterogeneous wireless technologies.

In order to characterize this downtime, we define the following experiment. We use the Iperf [14] tool to generate a UDP flow between the endpoints present at nodes labeled as N5 and N3 in Figure 2. After an instant, we switch off the mmwave interface of N5 connecting with N3 using the WIA. This action generates an OFPT\_PORT\_STATUS message at N5, which triggers the recovery procedure at the WiseHAUL SDN controller. Then, the WiseHAUL SDN controller reconfigures the FA of the different nodes with the appropriate OFPT\_FLOW\_MOD rules to resume the affected flow through the alternative path composed by nodes N5-N4-N2-N3. After some instant, we use the WIA to switch on again the mmwave

interface of node N5. A new OFPT\_PORT\_STATUS message is generated. This message triggers the recovery procedure, which reestablishes the former path, i.e, N5-N3. We define the downtime as the time interval when there has been an interruption in the reception of data packets at the endpoint of node N3 as a consequence of the switch off/on action. We repeat the same experiment 20 times with a mmwave link and a Wi-Fi link at the 5GHz band. Experiments were performed at night to reduce possible interference effects in the Wi-Fi link.

Figure 3 presents the experienced downtime when considering the different transitions and the different available wireless technologies. Figure 3 shows that the mean measured downtime after a switch OFF event is of 150ms for the mwWave interface, and slightly higher for the Wi-Fi interface, around 180ms. The packet reception resumes as



Paper Session X

Figure 3: Experienced downtime at different kind of wireless interfaces

soon as the FAs install the associated flow rules at the corresponding WXFЕs, as all the remaining interfaces in the alternative path are in ON state. However, the case for the OFF/ON transition is different. As we can see in Figure 3, the "mmwave downtime" is less than half the value (around 56 ms on average) with respect to the ON/OFF transition, and it is much lower than the Wi-Fi case, which is in the order of seconds. From the point of view of control operations requested by the WiseHAUL SDN controller to reconfigure the forwarding behavior of the WXFЕs, a lower downtime for the OFF/ON is expected because the recovery procedure implies more flow deletion operations in the different nodes (N5-N4-N2-N3) and less creation operations (N5-N3), which are more time consuming. In this case, the difference in experienced downtime comes due to the time it takes to the wireless device to make effective the order sent by the WIA. The mmwave link is established automatically after the interface is

ON again. However, in the Wi-Fi case, it is required to establish again the IBSS network forming the link. Since the information of previous link (frequency, bssid) is lost after the OFF operation, the creation of the new Wi-Fi link incurs bigger delays, around 5.2s on average. In the best cases, this delay is of around two seconds, however, in the worst cases, it could arrive to more than eight seconds.

## 5.2 Case Study II: Setting different Modulation and Coding Scheme

The following results show a TCP throughput characterization of the deployed mmwave network when setting the different available MCSs with the WIA under single and multi-hop communications and considering different maximum transmission units (MTUs). Differently to [5], we consider TCP connections between endpoints. TCP is selected given its significant share on Internet traffic. Additionally, this evaluation gains relevance for our mmwave setup, since the performance of TCP connections can be substantially affected due to the nature of such communications, i.e, high data rates transmitted over the different links of a multi-hop communication with potential channel variations. The evaluated flows last for ten seconds in each repetition, and are generated using the Iperf [14] tool. Every conducted experiment is repeated 10 times. For the case of one hop, TCP traffic has been generated between the endpoints of the nodes labelled as N2 and N3. For the case of two hops, traffic has been generated between the endpoints of the nodes labelled as N2 and N5, using N3 as forwarding node. TensorCom TC60G-USB3.0 EVK devices at N3 are configured as PCPs while devices at N2 and N5 are configured as stations. Table 2 present round-trip time (RTT) measurements between considered endpoints. These measurements have been obtained with 100 samples over 100 seconds.

Additionally, this characterization also explores the effect of the MTU on the throughput network performance. The maximum value of supported MTU of TC60G-USB3.0 EVK devices is 7912 bytes. However, we also consider 1500 bytes, which is popularly used in Wi-Fi networks and in Internet. As commented before, WiseHAUL platform combine different wireless technologies. This heterogeneity will be the norm in 5G network deployments and it is interesting to see how this can impact the performance of the network, since

the full capabilities of mmwave hardware would not always be possible to be exploited.

Table 2: Measured RTT (milliseconds)

Path	Minimum	Average	Maximum	Std Dev
N2-N3	0.522	0.647	3.04	0.127
N2-N3-N5	0.990	1.278	4.87	0.341

In our setup, we have considered that all mmwave links operate in channel 2 (60.48GHz), differently to [5], which considered different channels for the different mmwave links. As a member of the IEEE 802.11 family, the multi-hop communication in the presented setup is likely to be affected by co-channel interference (CCI), because both links share the same channel without being aware that they are in each other carrier sensing range. In addition to this, the CCI is also likely to occur due to the nature of the mmwave equipment under evaluation. As explained in [13], pencil-shape beams in mmwave networks are a myth. The reason for this is that real-world phased array antennas suffer from significant side lobes, and its effect is particularly noticeable in consumergrade hardware due to its cost-efficient design. Finally, it is worth mentioning that the IEEE 802.11ad standard only defines four channels, reducing the possibilities at designing the frequency planning and making co-channel interference more likely for close devices.

Figure 4 shows the performance comparison of end-to-end application throughput when considering the variables previously mentioned: available MCS, number of hops and different MTU values. As expected, the maximum throughput has been achieved using MTU 7912 bytes with one hop and using the fastest available MCS (MCS7). This value is of 843 Mbits/s on average. Notice that this is TCP throughput

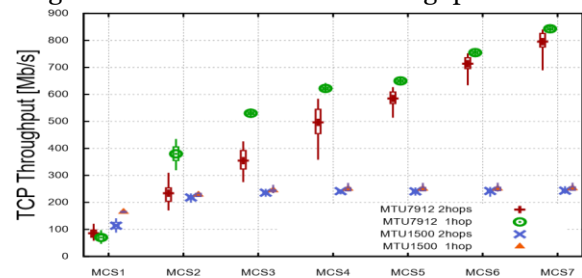


Figure 4: mmwave data plane performance characterization

and not link layer rate. When comparing the same MCS but different number of hops, we can see that the value

is very similar, around 800Mbps/s. However, the boxplots are wider (and in general for all the considered MCSs in the 2-hops communications), showing more performance variability due to possible fluctuations of link quality at each hop of the multi-hop communication. Furthermore, we can see that the setup does not suffer from CCI introduced by radiated side lobes. As indicated [13], the cause of a good performance at the receiver of the second hop is not only explained due to the selection of a good beam pattern capturing as most as possible of the main lobe of the desired transmitter but also due to very residual reception of side lobes transmissions from the other interfering transmitter.

If we refer to the values obtained with MTU 1500 bytes, we can see that using higher MCS schemes does not provide further gains in terms of end-to-end throughput and there is a saturation as depicted in Figure 4, and also experienced with TCP connections in [12], but for the case of one hop. In such cases, having network elements with a value of MTU equal to 1500 bytes, it could be convenient to use a lower MCS in mmwave devices (giving a similar performance than higher MCS) because they provide more reliable communications requiring less sensitivity at the receiver side, and at the same time, providing energy consumption gains due to reduced processing costs at the transmitter/receiver units.

## 6 CONCLUSIONS

In this paper, we present an agent (WIA) to manage the wireless interfaces of a real SDN-based multi-hop transport network research platform featuring mmwave and Wi-Fi communications. The experimental evaluations carried out suggest that such a RESTConf-based management agent is fully functional and enables fast prototyping, as opposed to more cumbersome options (e.g., modification of OF protocol). We employ the WIA to perform a quantitative evaluation of the deployed mmwave network, where downtimes in the order of 10s-to-100s of millisecond can be obtained in the presence of link up/down events. Furthermore, the maximum TCP throughput can be achieved in mmwave multi-hop scenarios despite neighboring links being configured in the same channel and using cost-effective pre-commercial equipment. Another relevant conclusion is that MTU mismatch among technologies in heterogeneous networks may hinder the full

exploitation of the capacity offered by mmwave links, which achieves its maximum throughput with big MTUs.

## REFERENCES

- [1] 2018. *Ath10k Candela Technologies CT 10.1 Firmware*. <http://www.candelatech.com/ath10k-10.1.php>.
- [2] 2018. *EU H2020 5G-Transformer, 5G Mobile Transport Platform for Verticals*. <http://5g-transformer.eu/>.
- [3] R. Santos et al. 2017. Small Cell Wireless Backhaul Reconfiguration using Software-Defined Networking. In *Procs. of the IEEE Wireless Comm. Networking Conf. (WCNC)*. IEEE, San Francisco, CA, USA.
- [4] J. Baranda et al. 2018. Orchestration of end-to-end network services in the 5G-Crosshaul multi-domain multi-technology transport network. *IEEE Comm. Mag. Net. and Service Man. Series 57* (July 2018).
- [5] R. Santos et al. 2017. Turning the Knobs on OpenFlow-Based Resiliency in mmWave Small Cell Meshed Networks. In *IEEE Procs of the 5GTestbed workshop, Globecom Conf.* IEEE, Singapore, Singapore.
- [6] Y. Zhu et al. 2014. Demystifying 60GHz outdoor picocells. In *Procs of the 20th Int. Conf. on Mobile Computing and Networking (Mobicom)*. ACM, Maui, Hawaii, USA.
- [7] Open Networking Foundation. 2012. *OpenFlow Switch Specification Version 1.3.0*. Technical Report ONF TS-006.
- [8] Open Networking Foundation. 2017. *4th Wireless Transport SDN Proof of Concept White Paper*. Technical Report.
- [9] IEEE. 2012. *Standard for Information technology—Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band*. IEEE SA. Technical Report.
- [10] J. Nunez, J. Baranda, I. Pascual, and J. Mangues-Bafalluy. 2016. Wise-HAUL: An SDN-empowered Wireless Small Cell Backhaul testbed. In *Procs of the 17th IEEE Int. Symp. on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. IEEE, Coimbra, Portugal.
- [11] S. Jumar, V. Vijay, A. Garg, and D. Koutsonikolas. 2016. Multi-Gigabit Indoor WLANs: Looking Beyond 2.4/5GHz. In *Procs of the IEEE Int. Conf. on Comms (ICC)*. IEEE, Kuala Lumpur, Malaysia.
- [12] K. Nguyen, M. Golam, K. Ishizu, and F. Kojima. 2017. Empirical Investigation of IEEE 802.11ad Network. In *Procs of the 10th IEEE Int. Workshop on Evolutional Techn. and Ecosystems for 5G and Beyond (WDN-5G ICC)*. IEEE, Paris, France.
- [13] D. Steinmetzer, A. Loch, A. García-García, J. Widmer, and M. Hollic. 2017. Mitigating Lateral Interference: Adaptive Beam Switching for Robust Millimeter-Wave Networks. In *Procs of the 1st ACM Workshop on Millimeter-Wave Networks and Sensing Systems (mmNets17)*. ACM, Snowbird, Utah, USA.

- [14] A. Tirumala, F. Qin, J. Dugan, and K. Gibbs. 2005. *Iperf: The tcp/udp bandwidth measurement tool*. <http://dast.nlanr.net/Projects>.
- [15] W. Xia, Y. Wen, C. Heng, D. Niyato, and H. Xie. 2015. A survey on Software-Defined Networking. *IEEE Communications Surveys and Tutorials* 17, 1 (June 2015), 27–51.